

SELF-ORGANIZED AIR TASKING:
EXAMINING A NON-HIERARCHICAL MODEL FOR JOINT AIR
OPERATIONS

John Bordeaux, Ph.D.

SRA, International
4300 Fair Lakes Court
Fairfax, VA 22033
(W) 703 803 1500
(F) 703 322 4965
john_bordeaux@sra.com

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE JUN 2004		2. REPORT TYPE		3. DATES COVERED 00-00-2004 to 00-00-2004	
4. TITLE AND SUBTITLE Self-Organized Air Tasking: Examining a Non-Hierarchical Model for Joint Air Operations				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) SRA International,4300 Fair Lakes Court,Fairfax,VA,22033				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 101	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Abstract

The hypothesis of this paper proposes that it is possible to structure a non-hierarchical approach to air tasking in the conduct of Joint air operations. For the private sector, advances in information and communication technologies have led to innovations in organizational structures in order to know more across the enterprise. However, the application of these “value network” principles has not been fully applied to the processes upon which the U.S. organizes for Joint force operations. A non-hierarchical model is constructed for the tasking of air assets in order to test an agent-based approach to the servicing of targets in an air campaign, using agent-based simulation techniques and models established by Epstein & Axtell (SugarScape) within the Santa Fe Institute’s Swarm agent modeling environment.

TABLE OF CONTENTS

	Page
1. Self-Organized Air Tasking.....	1
Hierarchies and Markets	1
Hierarchy and Air Tasking.....	4
A Non-Hierarchical Air Campaign.....	5
Motivation.....	8
Rationale.....	9
Hypothesis	11
Agent-Based Modeling and Complex Adaptive Systems	12
Potential Applications.....	14
2. Survey and Review of the Literature of Prior Research	16
The Evolution of Centralized Air Planning	16
Turning Point for Air Doctrine - 1943	16
Air Force as a Separate Service.....	18
Korea	20
Linebacker and Teaball	22
Operation Desert Storm and Instant Thunder.....	23
Operation Enduring Freedom.....	28
The Centralized Air Campaign	31
JFACC and DARPA.....	36
Agent-Based Modeling.....	37
Uses for Agent-Based Modeling	39
Architectures and Methods for Agent-Based Modeling	40
Sample Agent-Based Models	43
Economics: Trade Network Game Laboratory	43
Health Care: CoDiT and the SMASH Project (Martín, et al)	43
Intelligent Manufacturing: National Institute for Standards and Technology (NIST) / Advanced Technology Program (ATP)	44
Applicability of Agent-Based Models to Social Science Problems	44
Air Force Initiatives in Agent-Based Modeling	46
Continuous Air Tasking.....	46
Insect Pheromones	47
Agent-Based Modeling and Behavioral Representation (AMBR) Project.....	48
3. Approach.....	49
Research Approach.....	49
Research Designs & Methods.....	50
Research Method.....	50
Analytical Methods and Models - SWARM.....	51
Research Design - The Sugarscape Agent-Based Modeling Environment.....	52
Functional Description: Sugarscape	52
Modification of Sugarscape for this Work	54
ABACUS	57
4. Results	62

Lower Pk.....	64
Lower Fuel.....	66
Prefer SAMs	67
5. Summary.....	71
Bibliography.....	74
Appendix A - Annotated ABACUS Configuration File	78
Curriculum Vitae	82

LIST OF TABLES

<i>Number</i>	<i>Page</i>
Table 4-1 Baseline - Weapons Expended	62
Table 4-2 Weapons Effectiveness Halved – Expended Weapons.....	65
Table 4-3 Half Fuel Run - Weapons Expended	67
Table 4-4 Prefer SAMs - Weapons Expended.....	68

LIST OF FIGURES

<i>Number</i>	<i>Page</i>
Figure 4-1 Baseline Run – Targets Killed Over Time	62
Figure 4-2 Weapons Effectiveness Halved – Targets Killed Over Time	65
Figure 4-3 Half-Fuel Run – Targets Killed Over Time	67
Figure 4-4 Prefer SAMs – Targets Killed Over Time	68
Figure 4-5 Comparing SAMs targeted between Baseline and Prefer SAMs run.....	69

1. SELF-ORGANIZED AIR TASKING

Problem Statement, Motivation, Rationale, Hypothesis, Potential Applications

Hierarchies and Markets

The hierarchy is an organizational structure which represents a coping response to the high transaction costs associated with routine business functional interactions, which otherwise would require contracts for each interaction (Coase 1937). Given that common data abstraction and ubiquitous communication technologies are reducing transaction costs, some feel the hierarchy is a doomed organizational structure (Davidow & Malone 1992). However, the transaction costs associated with the growth of hierarchies are not only those that can be mitigated by improved communication flows; the cultural impediments of uneven trust relationships require a negotiation of intent prior to interaction in a non-hierarchical organization. Fukuyama (1999) shows that hierarchy is an established organizing principle, and the elements that lead humans to organize in this way are not obliterated by the information ‘revolution.’

According to Coase, the firm exists because, and only so long as, transactions within the enterprise are less expensive than transactions outside (or with outside entities). At the time he wrote, subcontractors often operated within a firm’s plant – departments negotiated each interaction with other departments by means of a contract. Coase showed that the hierarchy is a natural organizational structure for the firm, since it reduces the costs for interactions throughout the enterprise. By using the transaction costs as the measure of a firm, Coase provides insight into the “equilibrium of the firm,” applying the economic concept of marginal return to the organization.¹ By combining the ideas that the firm is organized and subject to an equilibrium regarding internal and external transactions costs, and that the firm is organized and managed by an entrepreneur who must decide in the face of uncertainty, Coase provides a basis for an inclusive definition of the “firm,” and a framework for future analysis of organizational theories.

¹ If the cost of management internal transactions exceeds the market cost of those transactions, either the firm will engage in market transactions for the function, or it will be overtaken by smaller firms who can execute the transaction for less cost.

As regards the firm, recent literature indicates the partial replacement of Coase's hierarchical organizational structures as the transaction costs of interacting with external functions decreases. Davidow and Malone assert that the "virtual enterprise will appear less a discrete enterprise and more an ever-varying cluster of common activities in the midst of a vast fabric of relationships (7)." "Much of middle management's function has been to serve as an information channel through which top managers can view events and to relay orders down to the individuals doing the work. These functions have become unnecessary because computer networks can carry much of the information about the status of operations more efficiently and effectively than can people (163)." They opine that ubiquitous information technology and the rise of the service industries will lead to a dramatic restructuring of existing business structures.

Quinn joins Davidow and Malone in claiming that hierarchies are "destroyed" and companies become "infinitely flat" as information technologies enable expanded spans of control. With this, Quinn dismisses overhead functions as "merely services the company has chosen to produce internally. Instead of blindly building self-owned, integrated services, companies are much more effective if they carefully benchmark, reengineer, and seriously consider outsourcing those where they are not best in [the] world. When this isn't done, corporate staff or overhead activities slowly and relentlessly tend to build into major bureaucracies (89)."

Where Fukuyama (1999) made the case for the enduring place of hierarchies in business organizations due to sociological imperatives, Davidow and Malone claim "the virtual corporation may exist in a state of perpetual transformation... [and] may appear amorphous and in perpetual flux, but it will be permanently nestled within a tight network or relationships (142)." The 'learning networks' have their theoretical foundation in complexity science (c.f. Waldrop), but are more appropriately understood as market-based organizational structures, where the hierarchy is challenged at each function and level to prove its contribution to the firm's competitive edge. This 'new' organizational structure is driven by customer relationships and the firm's core competencies in serving a dynamic marketplace.

Davidow and Malone write: “As the rapid gathering, manipulating, and sharing of information become a preeminent process and as company boundaries grow increasingly fluid and permeable, established notions of what is inside or outside a corporation become problematic, even irrelevant (140).” Here the authors continue their metaphor of dynamic, fluctuating firms interacting in a marketplace redefined as a network of evolving relationships, rather than existing in established roles of competitor, partner, or supplier.

A more complete argument for the elimination of hierarchies is found in Chapter Eight of Davidow and Malone, where the value of time is compared to a hierarchical structure of management. “...levels of management mean levels of approval, and levels of approval take time. The approvers become divorced from the market. Time is the virtual corporation's most valuable resource and the one commodity it cannot afford to waste (167).” A major force behind this restructuring of the enterprise is the emerging economies of speed. As compared to the economies of scale, which favored large firms with a robust infrastructure and redundant inventories - the factors associated with economies of speed include flexibility, rapid change of focus, and assessment and satisfaction of customer needs. The authors claim that this change in customer expectations (quick service and customer control of the product) means that speed will become the prime factor in the marketplace. Penzias claims that we have moved from this economy of speed to economies of convenience. These are cumulative eras, according to Penzias - economies of quality, speed, and convenience. Where most firms have achieved (or are achieving) speed, convenience is fast becoming the deciding factor for buying decisions for service products.

Davidow and Malone draw extensively from the work of Womack, Jones and Roos to show how the supplier-producer relationship has been redefined due to the Japanese model involving “kaizen.” This term, loosely defined as “continuous improvement,” characterizes the incremental but constant efforts to make business processes more efficient. “[The Japanese] have approached the process of improvement in a very conservative way, becoming great believers in taking small, incremental steps and pursuing the goal relentlessly over extended periods of time (127).” As these improvements are made within

an enterprise, firms are able to monitor more effectively the performance of their vendors, and begin to demand more efficient delivery of goods and services throughout the supply chain.

Hierarchy and Air Tasking

For U.S. military operations, the preferred organizing principle is the hierarchy. The orchestration of airpower, with its synchronization of target intelligence, aircraft weaponeering,² launching and recovery operations, flight paths, fighter escorts to engage enemy aircraft, positioning of air-to-air refueling resources – is no exception to this preference, although the centralized command and control of air operations is a relatively recent development in the history of U.S. military air operations. Driven by past errors which have been traced to a lack of central control, the Air Force has largely succeeded in crafting a hierarchy for air operations: the Joint Force Air Component Commander (JFACC), which has operational control of all U.S. aircraft engaged in a military conflict, regardless of Service affiliation. This is heralded as an improvement to the effective and efficient use of airpower, and was tested on a large scale for the first time in 1991's Operation Desert Storm, also referred to in the literature as the Gulf War. Winnefeld and Johnson found that: "The mode that has worked best to date, as confirmed by the experience of the Gulf War, is for one component commander to act as the lead commander and be given tactical control of sorties from the committed assets of the other services. This functional air component commander should have a joint staff and senior representatives of the coordinated components' forces on duty as his air operations center, and in some cases those representatives should be at the deputy functional component commander level (150)." While naval aviation forces are also subject to the JFACC, and have developed operational concepts for a "JFACC Afloat," these air assets have historically functioned in a decentralized manner.

² The loading of a specific weapon onto a specific aircraft. Certain aircraft are optimized for the delivery of certain weapons, and those weapons in turn have varying degrees of effectiveness against certain types of targets. The careful planning of which type of weapon, aboard which plane, will hit which target, is the focus of the Master Air Attack Plan.

A Non-Hierarchical Air Campaign

One question that arises from this review is: How does one reconcile the insights of organizational theory, and its contributions to the modifications or elimination of the hierarchy, with the movement towards increased centralization for U.S. military air operations? Might the increased effectiveness and efficiencies noted earlier for “learning network” organizations be applied within the context of U.S. military air tasking? If there were a way to retain centralized command of U.S. airpower while introducing a decentralized execution, is the current approach to resource allocation optimal for operating within this paradigm? It may be that allowing for a decentralized control of airpower – allowing aircraft to use advanced information and communication technologies to (possibly) more effectively and efficiently engage targets, in effect, to self-organize to achieve objectives – is one way in which the benefits from these technologies can be realized within the context of U.S. military operations.

This research is an existence proof to study the feasibility of a self-organizing air campaign³ for Joint air operations. Using agent-based modeling techniques, a simulation is constructed to test the concept of instituting an self-organizing system – wherein constructs of pilots, aircraft, weapons, escorts (all acting as agents) – bid on fixed/known as well as mobile/emergent targets in a Joint air campaign. By specifying the “best” weapon for a particular target, and allowing aircraft to communicate their capabilities and position to nearby aircraft, the objective is accomplished of pre-planned aircraft tasking: allowing the most appropriate aircraft to attack the target. This objective is achieved without a pre-planned air tasking order (ATO), which specifies before the day’s activities which aircraft will strike which target.

In the years since air power was first used to achieve war aims, each conflict has provided insights into its effective use. The application of these insights provides for the evolution of the “air campaign” concept. Following Operation Desert Storm (ODS), two such lessons provide the key issues for this paper:

³ This term is used throughout the paper, and refers to the planning and execution of airpower used in support of objectives in a military conflict.

A scripted⁴ air campaign that treats the enemy infrastructure as a system (by using parallel attacks against targets that are critical nodes in the enemy sub-systems), can be extremely effective in achieving operational and strategic objectives. (Warden, others)

Air attack assets will become unexpectedly available, and targets will emerge throughout the campaign, “inside the planning cycle” – often before the planning process can react and re-assign those assets to those targets. A completely scripted air campaign, as effective as it is, by definition does not address these contingencies.

These lessons were predictable, as Warden, the chief architect of the ODS air campaign, wrote in 1989: “In the process of planning or executing an air campaign, three especially thorny issues confront the commander and his planners. The first is the use of air in emergency situations, such as a fast-progressing enemy ground offensive; the second is deciding on the relative effort to be assigned interdiction and close air support; the third is the desirability of carrying out air superiority, interdiction, and close air support simultaneously (Warden, 132).” While the air planners in ODS were able to adapt to emerging targets and unassigned assets through ad hoc measures, peacetime research efforts should attempt to incorporate the lessons of previous conflict in order to help evolve the air campaign concept. This research aims to address the re-tasking of air assets when events occur outside of the main air tasking order (ATO) “script.” This “script” is formally referred to as the Master Air Attack Plan (MAAP).

The Master Air Attack Plan begins the air planning process. Inputs to this process include the JFACC daily guidance, unit level resource information, close air support⁵ (CAS) and defensive counterair⁶ (DCA) sorties, airspace control plans, target nomination list (TNL), intelligence data, current and predicted theater wide weather, other component plans

⁴ Refers to the planning process that details aircraft, weapons suite, time of take off, primary and secondary targets, and other elements amounting to a “script” compiled by planners, and followed by the pilot during the air attack.

⁵ Aircraft assigned to attack opposing ground forces in proximity to and presenting a danger to friendly forces.

⁶ Aircraft assigned to engage and destroy enemy air assets. Defensive counterair refers to missions designed to reduce the enemy’s ability use aircraft, and includes strikes against runways or aircraft maintenance facilities. Offensive counterair refers to the engagement of enemy aircraft aloft. Aircraft are usually designed for a specific role, e.g., the F-15C is designed for offensive counterair, while the A-10 carries air to surface munitions and is well suited to the ground attack tasks associated with defensive counterair missions. The F-15E and F-18 are examples of multi-role aircraft, and feature the maneuverability of the air-to-air aircraft while carrying air to surface munitions.

(ground, air, maritime), and the Joint Munitions Effectiveness Manual (JMEM). Planners review the candidate targets for the next ATO, and group the targets based on either geographical location in the case of a large theater, or mission requirements (which sorties require an escort jamming aircraft, for example). The planners then determine the available sorties based on knowing how many of each aircraft type are available, and identify support requirements (jamming⁷, air refueling, forward air control⁸). The outputs from the MAAAP process are a revised airspace control plan, electronic combat plans, and an airspace control measure request.

The Air Tasking Order is a computerized daily list of air assets, broken down by missions, squadrons assigned, targets, restricted operating zones, low-level transit routes, drop/landing/extraction zones, and air refueling areas⁹. It does not specify tactics or flight plans, but does include basic mission information: mission number, priority, mission type, time on and off target, alert status, location, call sign, number and type of aircraft, ordnance type, IFF/SIF (identification, friend or foe/selective identification feature) mode and code¹⁰, and time and target location (Sessions & Jones 1993). The ATO builds on the airspace control plan and electronic combat plan, incorporating the approved JFACC daily guidance to develop the final plan for the air campaign over the next 24-48 hours. The scope is likewise broad. For the Balkans Combined Air Operations Center, a staff of 300 work to produce the ATO for their area of responsibility. Originally established to support

⁷ Electronic combat includes the exploitation of the radioelectromagnetic spectrum in order to deny the enemy the ability to communicate effectively with its forces. Jamming refers to the broadcast of noise across a radio or telemetry circuit, in order to render the affected frequencies unusable.

⁸ Air control is similar to the function served by air control towers in civilian airports. In conflict, the air control function is usually assigned to an airborne platform, such as the Airborne Warning and Control System (AWACS). Often, this air control is delegated to friendly ground forces near the target area (*forward air controllers*); forces specially outfitted with communication gear to assist in guiding friendly aircraft safely into and out of the target environment.

⁹ These terms refer to various elements in orchestrating the take offs, landings, transit and activity of aircraft. Because the U.S. is attempting to operate aircraft in an unwelcome manner, the basic elements of airspace control provided by civilian airports are lacking, or at least unavailable to friendly forces. In addition, the aircraft in question are making sudden changes in altitude, and are coming under fire. Finally, some of the aircraft are engaged in destroying what civilian air control infrastructure does exist in the target country. Much of the air planning, then, is aimed at establishing a dynamic air control environment to minimize risk to friendly aircrews and maximize the probability of success in the air campaign.

¹⁰ Essential for questioning aircraft in flight to determine friendly or enemy status, the IFF/SIF method is a challenge-response way to authenticate that an aircraft is friendly.

Operation Deny Flight (enforcing UN sanctions over Iraq), the Balkan center controls aircraft from bases throughout Europe. (Caddell).

In examining the feasibility of a self-organizing air campaign, this paper implies that the partial decentralization of tasking for air assets may also be feasible. Many historians will question this feasibility, particularly following the experiences in the Vietnam conflict, which predate the establishment of Joint air planning. In Vietnam, excessive geographic decentralization of air planning often meant the lack of coordination among Services or areas of responsibility (AORs) within the theater of operations. With no single air commander, with authority over Air Force and Navy fixed-wing operations or Army helicopter missions, it was difficult to manage assets for strategic theater-wide effects. Also, targets could appear on the air tasking lists of more than one Service, resulting in redundant capability expended on some targets. With the advent of improved information and communication technologies, decentralization may no longer mean a lack of coordination. Information technology has a democratizing effect, where the marginal transaction cost (in terms of money, time, and effort) of delivering information is sufficiently low as to remove barriers to long-distance, real-time coordination mechanisms.

Motivation

“Information to the warfighter.” This call to arms characterizes much of battlefield command and control (C2) efforts since Desert Storm. We have “intelligent” weapons systems, with “smart” guidance packages. We tell weapons where to detonate, and loose them at targets. “Fire and forget” is a term often used among air warriors to describe these techniques used to deploy “smart” bombs and missiles. The pilot, however, must still work according to a script worked out hours earlier. We are still using air assets as we did fifty years ago – manned aircraft working to a script, and releasing ordnance over a target as described in the ATO. As we increase the information available to the warfighter, perhaps new ways of assigning weapons to a target can be found. The recent campaign in Afghanistan is a departure from the scripted approach, but the scarcity of the target set may not be representative of future theaters. Likewise, while the use of patrolling bombers and unmanned combat air vehicles (Predators with Hellfire missiles) represent new uses of

existing air assets, it is too early to conclude that the U.S. has abandoned the centralized, scripted approach to an air campaign.

In business, the increase of available information, when coupled with a change in process or workflow, may lead to increased efficiencies and effectiveness of individual agents in an enterprise. Likewise, we can envision situations where the increase in information may result in inefficiencies where workflow or processes are not changed to accommodate the glut of information. In complex adaptive systems, the agents who can adapt their approach to a changing environment help sustain the system at “the edge of chaos” – the most lucrative formation for the achievement of system objectives. For the warfighter, then, we should investigate new processes, ones that allow adaptation to a changing environment, in hopes of discovering new efficiencies and increasing the effectiveness of warrior assets.

Rationale

As part of legislation establishing the Quadrennial Defense Review, a Congressionally-mandated review of force structure and mission for the Department of Defense, The National Defense Panel (NDP) was also established. This civilian board has a charter to review the QDR results by providing a critique of the Secretary of Defense’s conclusions following each review. For 1997, the NDP report included a recommendation that the DoD “exploit information technology to integrate forces and platforms more effectively (National Defense Panel).” With the common objective of providing more and faster information to the decision-makers/stakeholders in a conflict, the deployment of technologies alone will not be sufficient to realize the potential benefits. Builder, et al., provide a strong case for the development of a concept for command and control that maximizes (or reaches beyond) available technology – rather than treating the command and control processes as simple implementations. Given the process and organizational innovations in the private sector following the incorporation of advanced information and communication technologies, it behooves the researcher in the public sector to consider similar innovations for processes within the public sector.

This paper seeks to initiate a body of work that may lead to a more effective integration of weapons platforms (manned and unmanned aircraft) in Joint air operations. In many ways,

the lessons of Operation Desert Storm (ODS) were optimistic, in that an enemy presented the Coalition forces with months of preparation time, the ability to therefore marshal half a million U.S. troops and thousands of aircraft, and time to plan a devastatingly synchronized attack to its ground forces and strategic centers of gravity. Winnefeld, et al. refer to the Persian Gulf as a “near-ideal theater for deploying a large expeditionary force (26).” Future conflicts may force air planners to work with fewer weapons platforms. When an attack sortie cannot hit its target due to weather or other factors, the best use for that aircraft is to strike a secondary target, and not return to base with its ordnance intact. Luttwak summarizes the reasons for failure to strike primary targets in ODS: “Even against targets of known location, bad weather seriously impeded the air attack – partly because the most accurate precision bombs must still be guided down visually and partly because the effect of prior sorties could not be photographed to establish if further attack was warranted (23).”

Given that the U.S. military is investing heavily in information technologies in order to provide heightened visibility into the battlefield situation (enemy and friendly force capabilities, vulnerabilities, intention, and location); the persistence of a scripted ATO that relies on 24-72 hours-old information to make targeting decisions is anachronistic. We seek here to apply processes to air targeting that leverage the future battlefield, by allowing information about the battle to be immediately incorporated into the decisions regarding which aircraft strike what targets.

The most recent example of the use of airpower to accomplish national objectives is the campaign in Afghanistan. The ratio of precision weapons versus unguided ones, combined with the use of Special Forces in place guiding the bombers to their destination, may seem to make this work somewhat obsolete. Rather than a 72-hour planning cycle, U.S. bombers were able to loiter in the area, directed to targets of opportunity as they were spotted by U.S. ground forces. The experience from the Afghanistan campaign, however, must be tempered by the fact that a sparse and largely undefended target set, combined with friendly forces throughout the country, provided an environment that should be considered hospitable to the successful prosecution of an air campaign. It would be premature to assume with confidence that, in a future conflict, the U.S. always would gain a rapid

victory over air defenses, leverage large friendly forces in country, or enjoy near-complete freedom of movement on the ground for U.S. Special Forces to aid in air targeting. If the U.S. brings airpower to bear on an enemy force against a more defended nation-state (e.g., Iraq, Iran, North Korea); it will likely face an integrated air defense system and more heavily defended high-value targets. The centralized planning process has not been changed, and the approach to a conflict of this scale would more likely resemble an Operation Desert Storm rather than Operation Enduring Freedom's campaign in Afghanistan.

Hypothesis

Current organizational structures for Joint air operations are potentially inefficient, when considered within the context of the possibilities offered by advanced information and communications technologies. These advances in technology are accompanied by innovations in organizational structures, at least in the private sector, in order to know more across the enterprise. The Department of Defense has studied these innovations and extended them across the acquisition community as best principles for supply chain integration and management. However, the application of "value network" principles has not yet been applied to the processes through which we organize for joint force operations. The hypothesis for this work can be stated thus:

It is possible to structure a feasible non-hierarchical approach to air tasking in the conduct of Joint air operations.

A simulation is constructed of a non-hierarchical, or self-organizing, assignment of targets within an air campaign, using agent-based modeling, as an existence proof. Allowing for the semi-autonomous interaction of aircraft/weapons/support packages as they bid for targets would represent a new approach to this information-intensive process, currently accomplished through manual planning (where each target is assigned to a specific aircraft with a defined weapons suite). The hypothesis is rejected if the self-organizing air tasking model is infeasible.

The degree of autonomy granted to agents is an issue for organizational theorists and sociologists. How 'necessary' is a directed hierarchy in an enterprise where each

stakeholder is committed to a common core of principles, understands the mission, and has the information necessary to maximize their effectiveness within the enterprise? The implied question here of ‘necessity’ overlooks the sociological view (Fukuyama 1999) that hierarchies do not exist solely on the basis of transaction cost economics; instead, there are indications that human beings engaged in a common endeavor will often prefer a hierarchical organization to complete agent-level autonomy.

This work, however, does not address the social feasibility of a partially decentralized approach to air tasking – we leave the political and human teaming variables to future research. The metrics of feasibility here are intentionally narrow, relating only to the achievement of operational objectives, as this represents the first step to investigating the potential promise of decentralized, self-organizing air campaigns. The hierarchical organizational structure may be reviewed in light of emerging information and communication technologies, but there are sound principles that present barriers to decentralization, principles that must be addressed in any non-hierarchical model. The key barrier for air operations in conflict is time-to-decide. Time is the key variable here, as the increased time to assign and decide in any market-based system for assignment of targets may reduce flexibility and effectiveness. Feasibility here, then, is assessed according to the following questions:

1. Are the targets struck quickly, that is, within a reasonable amount of time after becoming candidates to the target list?
2. Are the targets struck effectively in this construct – with sufficient capability, as needed to damage/destroy the target?

Metric 1 is related to the inherent inefficiencies noted in agent-based simulations of economic activity (Epstein & Axtell); while metric 2 captures the normative validation of the model: Does it resemble a reasonable use of airborne attack assets against targets?

Agent-Based Modeling and Complex Adaptive Systems

In suggesting a partial shift in the organization of air operations, to allow for the decentralized air tasking, the air campaign should be viewed as a self-organizing system,

wherein agents purposefully work to achieve local objectives. This is in keeping with the use of agent-based modeling to examine whether local rules can be written for semi-autonomous or autonomous agents that accurately reflect the design of a grand strategy. If we can enable agents to execute based on fluid local situations, and achieve the same adherence to the grand strategy as is found in strict centralized execution, we may realize a greater flexibility and possibly therefore increase the chances for success. In speaking at a conference entitled, “Embracing Complexity”, Fukuyama reflected on what this meant for the German army in 1940:

“[T]he Germans had a concept of leadership that is the basis of what the U.S. Army now calls ‘Mission Orders.’ Leaders were taught not to micromanage, but only to set down the most general kinds of objectives and to push as much authority as possible down to the lower levels of the organization. That’s what led to the victory over France, because every platoon commander and brigade commander in the German army was given a tremendous amount of leeway, based on commonly shared norms and values. It enabled them to innovate and take advantage of local opportunities. It created an adaptive system (Fukuyama & Haeckel).”

Viewing a military organization as a complex adaptive system, and reducing the amount of decisions reviewed at higher levels of command, is therefore not a new perspective.

While Deming’s work with Japanese firms in the 1950s was the first to recognize that commercial firms were behaving as complex adaptive systems, the Santa Fe Institute, academic home to work in the field of complexity science, first developed the large scale application of agent-based modeling techniques against complex adaptive systems when modeling the interactions within and among economies. The powerful message lay in their results for the “successful” simulations of complex systems: begin from the bottom up, because that is how life begins. Instead of engaging in aggregated simulations of a system, they modeled the component ‘agents’ within the system, and empowered them with rules for adaptation to their environment and interaction with other agents. One economic model briefly described by Waldrop contained 4500 equations and 6000 variables. “And yet none of the models really dealt with social and political factors, which were often the most important variables of all (93).”

When studying the self-organizing behavior of complex adaptive systems, the use of bottom-up, or agent-based modeling techniques, is preferred. The Argonne National Laboratory maintains a Complex Adaptive Systems program which develops agent-based modeling and simulations of physical, economic, and social systems.¹¹ Iowa State University's Department of Economics develops software for Agent-Based Computational Economics and Complex Adaptive Systems.¹² Cap Gemini Ernst Young's Center for Business Innovation found that "agent-based simulation is an especially promising complex adaptive systems technique that serves as a platform for emulating the nonlinear attributes of real-world complex systems. Among the dominant forms of modeling in the social sciences today is game theory, which is built on rational choice assumptions. Agent-based modeling is an alternative method most suitable when the system to be studied exhibits adaptive behavior."¹³

Potential Applications

Assuming an automated system is possible for this crucial planning task, one can imagine the implications for weapons acquisition informed by such simulations. For example, with a robust set of scenarios for analysis using an agent-based air campaign model, the analyst may gain insight into force structure implications. One hope here is that this could lead to an increased focus on providing the "right" capabilities against a target set, apart from the political pressures that may lead (unnecessarily?) to a multi-Service representation in most conflicts.

Simulation-based acquisition (SBA) – These efforts are now focused primarily on identifying architectures, common data catalogs, and simulations that show promise. In addition, SBA efforts are seen within programs (Crusader) as increasing visibility across products. This approach, if sound, may add to current SBA efforts as a way to explore an objective use of force capabilities, and may indeed lead to innovations in the use of force packages in certain scenarios.

¹¹ http://www.dis.anl.gov/msv/msv_cas.html

¹² <http://www.econ.iastate.edu/tesfatsi/acecode.htm>

¹³ <http://www.cbi.cgey.com/research/current-work/biology-and-business/complex-adaptive-systems-research.html>

Force structures – The current disagreements regarding appropriate force packages may be further informed by the market allocation of capabilities against target sets within a scenario. For example, as is the case with the use of air power against enemy ground forces – does an increased capability in guided “intelligent” munitions imply a reduction in the need for the symmetrical use of U.S. ground forces?

Remotely-Piloted Vehicles (RPV) – One promising development, as seen most recently in Afghanistan, is the use of RPVs¹⁴ to engage and kill targets is a promising development in the use of airpower in military conflict. The Predator RPV was used to extend the eyes of its ground-based crew, who steered the aircraft and fired missiles from the platform at potential targets. If the self-organizing air campaign is feasible, there are potential implications for expanding the use of RPVs. Instead of merely providing a remote intelligence collection and firing platform, the “intelligent” RPV may be also imbued with onboard logic that mirrors agent rules for behavior, bypassing the need for pilot judgment regarding target selection.

Multi-Service representation – Analysts question the use of land-based air forces in the prosecution of El Dorado Canyon (1986 bombing of Libyan air defense and infrastructure targets – see Winnefeld & Johnson) when carrier-based air forces were within range and in sufficient numbers. The Air Force lost an aircraft returning to a distant base in England. Was this loss necessary? Did the Air Force “need” to be employed against these targets, or should the range issues (and the refusal of the French to allow overflight rights that would have reduced the length of the routes) have left this as a Navy operation? A non-hierarchical simulation may lead to the analysis showing that such far-flung assets are not necessary in order to service such targets when suitable and proximate assets present less risk to U.S. pilots.

¹⁴ The alternate term found in the literature for these aircraft is Unmanned Aerial Vehicle (UAV). The term describing a concept wherein these platforms are outfitted with weapons, for example the Hellfire missiles used in Afghanistan, is Unmanned Combat Aerial Vehicle (UCAV).

2. SURVEY AND REVIEW OF THE LITERATURE OF PRIOR RESEARCH

Evolution of Centralized Air Planning, Agent-Based Simulation, Current Air Force Initiatives in Agent-Based Modeling and Simulation

The Evolution of Centralized Air Planning

The history of air warfare is necessarily brief, and dates from the Italo-Turkish war of 1912, where aircraft were used for combat purposes beyond simply extending the reconnaissance capabilities of the ground forces. For the purposes of this paper, we are concerned primarily with the salient historical points that led to the establishment of a directed hierarchy for the use of U.S. air forces in conflict.

During World War I, aircraft were used primarily for reconnaissance in order to gain intelligence regarding enemy troops movements. Douhet wrote regarding the potential of airpower in changing the way wars were fought, and remains one of the chief architects of air warfare, recognizing the revolutionary impact that control of the air can have on conventional combat. At first, air assets were used to support ground forces. The defining U.S. doctrinal document issued in 1926 specified that the “purpose of air units was to aid the ground forces by destroying enemy planes and attacking enemy ground forces (Basic Aerospace Doctrine A-1).” Perhaps more importantly, Army commanders controlled the air units, and therefore they alone decided the proper employment of air power.

Turning Point for Air Doctrine - 1943

Air Force doctrinal documents assert that 1943 represents a defining moment, separating the development of air doctrine, that of support to ground forces, from the modern doctrine, which adopts a measure of Douhet’s exhortations regarding the promise of aerial bombing. This doctrinal shift followed the experiences in North Africa, following Operation Torch. Air power enthusiasts believed that air power could be used to advance the war aims. Freed from having to support a fielded army, some felt that the “proper” use of air power

might obviate the need for ground forces in some conflicts. This position has been widely assailed, but following the NATO operations against Yugoslavia in 1998, one of the most vociferous opponents of this view, British historian John Keegan, was forced to admit that perhaps air power had finally “won” a conflict.

In Operation Torch, the opening of the North Africa campaign, the U.S. forces entered their “first large-scale ground-air operation” with the air assets subordinated to the ground force commanders. More than that, they were subject to the doctrine of the time, which stated that the principle mission for airpower was to support the ground forces – and that therefore they were subordinate to the ground force commander. The air commander operated under the Army commander, who could allocate aircraft to lower-level ground units. “FM 31-35 stated unambiguously that ‘the most important target at a particular time will usually be the target which constitutes the most serious threat to the supported ground forces. The final decision as to priority of targets rests with the commander of the supported unit.’ This was in keeping with the traditional ground-force view of aviation as a primarily defensive weapon. Such operations robbed airpower of its greatest strengths--mobility and flexibility--and made it impossible for the air units to achieve air superiority (Frisbee 1990).”

Since the command of air forces remained under the local ground force commander, there was no effort in Operation Torch to achieve air superiority or even plan for a coordinated use of air power in the theatre to provide effective ground force protection. This led to interest at high levels to review the operational use of air power in light of its potential strategic advantage. Following a 14 January 1943 meeting at Casablanca between Roosevelt and Churchill, Gen. Eisenhower reorganized the Allied forces in North Africa into component air, ground, and naval commands reporting to him. His air component, the Mediterranean Air Command, was headed by Air Chief Marshal Sir Arthur Tedder and included as its principal element Northwest Africa Air Forces, commanded by General Spaatz. Under Spaatz were separate strategic, tactical, and coastal air forces, service and training commands, and a reconnaissance wing... The air force was to be governed by doctrine developed by Air Vice Marshal Coningham and General Montgomery in the campaign against the Afrika Korps. That doctrine was laid out in a pamphlet which,

though signed by Montgomery, is thought to have been written by Coningham. It stated, in part:

‘The greatest asset of airpower is its flexibility.... So long as this is realized, then the whole weight of the available airpower can be used in selected areas in turn. This concentrated use of the air striking force is a battle winning factor of the first importance. It follows that control of the available airpower must be centralized and command must be exercised through Air Force channels...Nothing could be more fatal to successful results than to dissipate the air resources into small packets placed under command of army formation commanders, with each packet working on its own plan. The soldier must not expect, or wish, to exercise direct command over air striking forces.’ (Frisbee 1990).”

Air Force as a Separate Service

Following World War II, the Services engaged in a lengthy battle regarding the future organization of the U.S. military. While Army Air Force Gen. Spaatz testified before the Senate Committee on Military Affairs that “unity of direction [and] equality for the Air Force which will insure unification of our air potential [were] absolute imperatives which stem from the lessons of [World War II] (Futrell Vol. 1, 192),” the Navy in particular feared the consolidation of air assets in a single air force would compromise the integration of Marine aviation and naval assets, and subordinate Naval carrier aviation to this new Service.

Following months of debate, the National Security Act of 1947 established the Air Force as a separate arm of the U.S. military. This Act also established National Military Establishment, with components including the National Security Council, Central Intelligence Agency, and others. In addition to the Act itself, President Truman signed an executive order on the same day he signed the Act, naming James Forrestal as the first Secretary of Defense, but also further defining the role for the U.S. Air Force, charging it to: “organize, equip and train air forces for air operations including joint operations; to gain and maintain general air superiority; to establish local air superiority where and as required; to develop a strategic air force and conduct strategic air reconnaissance operations; to provide airlift and support for airborne operations; to furnish air support to land and naval forces including support of occupation forces; and to provide air transport for the armed forces except as provided by the Navy for its own use (Futrell Vol. I, 196).”

According to the diaries of Secretary Forrestal, the fears of the Navy regarding the attempt by the new U.S. Air Force to centralize air assets were well founded. “I remarked that there were these fundamental psychoses, both revolving around the use of air power: 1. The Navy belief, very firmly held and deeply rooted, that the Air Force wants to get control of *all* aviation; 2. The corresponding psychosis of the Air Force that the Navy is trying to encroach upon the strategic air prerogatives of the Air Force. (He denied that the Air Force had the first objective in mind, although that is what General Spaatz has said to me in private conversation) (Millis 466).”

When the dust had settled regarding the new organization, the Act allowed for the maintenance of four air forces: retaining the Marine tactical air forces as organic to their ground divisions; allowing for the continued development and deployment of carrier-based aviation; and for the assignment of rotary-wing aircraft within the U.S. Army. “The proverbial stone in the shoe was one sentence that appeared in the Army portion of both documents, section 205(E) of the National Security Act of 1947, and Section IV of the ‘Functions Papers.’ That statement read as follows: ‘The United States Army includes land combat and service forces and *such aviation* and water transport as may be organic therein (Tilford 16,17).” This agreement insured that the Army would retain operational control over some organic air assets, perpetuating a counterweight to the arguments for centralized control of air assets.

The remaining tensions continued, with the Air Force interested in consolidating air power under the single Service, and the other Services resisting a change to their organic capabilities. While preparing the budget for 1948, it became clear that Truman was interested in supplemental defense spending, but the Services were irresolute in how the additional monies should be spent. The Air Force wished to finance 70 combat air groups, while the Navy wanted additional carriers and carrier-based aircraft, the latter to be equipped for atomic weapons. The Air Force viewed this as an encroachment on their ownership of atomic weapons, and seized upon the opportunity to once again argue against the existence of Navy aviation assets outside the purview of the Air Force. Because of the lack of consensus, Secretary Forrestal convened an extraordinary meeting of the Joint Chiefs of Staff in Key West, Florida in March of 1948 to resolve these issues – vagaries

remaining from the legislation. The final statement (“Functions of the Armed Forces and the Joint Chiefs of Staff”) issued from this meeting recognized the following three relevant points of agreement (taken from Secretary Forrestal’s diary):

“For planning purposes, Marine Corps to be limited to four divisions with the inclusion of a sentence in the final document that the Marines are not to create another land army.

“Air Force recognizes right of Navy to proceed with the development of weapons the Navy considers essential to its function but with the proviso that the Navy will not develop a strategic air force, this function being reserved to the Air Force. However, the Navy in the carrying out of its function is to have the right to attack inland targets – for example, to reduce and neutralize airfields from which the enemy may be sorting to attack the Fleet.

“Air Force recognizes the right and need for the Navy to participate in an all-out air campaign (Millis 392-393).” The Air Force remained interested in a consolidation of air power, but was also convinced that atomic weapons would be the strategic weapon of choice in future conflicts and primary fought to gain and retain the mission of “strategic” air power – understood to refer to the delivery of atomic weapons. Following the meeting, General Spaatz balked at agreeing to the points as specified. Forrestal relates in his diary that Spaatz remained concerned at the duplication of effort that came about from the existence of multiple air forces, but that he acknowledged resolving that issue meant a change in the law (National Security Act 1947), and acquiesced when Forrestal reminded him that they were to implement the law, not change it (Millis 394).

Korea

“The Korean War was the first conflict to test the unified military forces of the United States. Although the U.S. Joint Chiefs of Staff had directed the Far East Command to provide itself with a joint command staff adequate to ensure that the joint commander was fully cognizant of the capabilities, limitations, and most effective utilization of all the forces under his command, the United Nations Command/Far East Command operated for the first two and one-half years of the Korean War without a joint headquarters. Practically all of the interservice problems which arose during the Korean War could be traced to misunderstandings which, in all likelihood, would never have arisen from the deliberations of a joint staff. In the absence of the joint headquarters staff, the full force of United Nations airpower was seldom effectively applied against hostile target systems in Korea. (Futrell 1983, 693).”

The onset of the Korean conflict in 1950 found the U.S. air forces unprepared for a robust defense of the Korean peninsula. The lack of a joint commander led to significant disconnects, notably for our purposes in the area of target selection. Kropf points to the leadership style of General MacArthur to explain the lack of coordination for joint operations, and the development of hostility among his senior staff: “MacArthur, heading the Southwest Pacific Command, surrounded himself with a staff of trustworthies (some say sycophants) known as the “Battan Gang” and kept his theater headquarters far from the front (1990).” Until the installment of General George C. Kenney as air component commander, MacArthur’s approach had left the air and naval component commanders subordinate to the ground component commander. The Air Force continued to pursue its doctrine of deep interdiction¹⁵ by moving bombers from Strategic Air Command to Far East Air Force (FEAF) control in Japan, but the leadership there was under pressure to support the Allied forces in contact with the enemy. In fact, the FEAF was unable to manage the target selection process, as the General Headquarters (GHQ) staff established the GHQ Target Group and directed air operations from Tokyo.

“The Target Group, made up of GHQ staff officers, ‘lacked the experience and depth of knowledge for targeting an air force. . . . [T]he [Target Group] effort was inadequate.’ As an [ex]ample, 20 percent of the first 220 targets designated were nonexistent, such as the rail bridges at Yongwol and Machari—two towns without railroads at all. A GHQ Target Selection Committee, which included high-level USAF and US Navy personnel, was formed to improve targeting. The GHQ Committee did improve performance but was dependent on the FEAF Formal Target Committee, with Navy, Fifth Air Force, and Far East Bomber Command representatives providing expert targeting. This FEAF Committee did not get full authority for air targeting until the summer of 1952, two years into the war. The overall effect was the failure to fully integrate air power into the theater campaign (Kropf).”

As a substitute for centralized joint operations, the Services initially settled on “coordination control” as the operational concept for air operations in the Korean theater. Meant as a way for the air arms of each Service to operate in concert with the objectives and requirements of the Far East Air Forces, it instead fell victim to definitional problems.

¹⁵ Refers to the targeting of supporting infrastructure elements to reduce the enemy’s ability to move and supply its forces; e.g., petroleum, oil, and lubricant (POL) stores, rail lines, transshipment points, stockpiles, etc.

Winnefeld & Johnson found that “[t]o Air Force staffs and commanders it was thought to confer the authority to designate air missions and task in Korea to be performed by the air components of all services. In short, it precluded the Navy from conducting its own air war in Korea apart from the Air Force (42).” Not surprisingly, the directive was interpreted somewhat differently from the Navy perspective. Naval staff believed “coordination control” to mean that the FAEF could request but not task naval air support from Navy resources. “The key question for both FEAF and NavFE [Forces Far East] was who tasked the Navy to provide sorties, air coverage, and so on...in effect the Navy was autonomous in deciding which and how many of its assets to commit to a given mission or task in support of the other two component commands of CINCFE [Commander-in-Chief, Far East] (43).” The lack of a single joint force air component commander with the authority to resolve this issue contributed to confusion and a reduced unity of command for the early years of the Korean conflict.

Linebacker and Teaball

During the Vietnam conflict, air operations were planned separately by each Service (Army, Navy, Air Force, and Marine) specific, and targets were likewise assigned within the tactical target lists.¹⁶ Coordination was minimal, as indicated in Tilford: “The Air Force log included examples of incidents where Army helicopters and Air Force T-28s strafed or bombed the same area at the same time without notifying one another (Tilford 69).” Murray lamented the absence of that “tired but still crucial principle of unity of command...the absence of a single air commander produced chaos (107).”

While the inter-Service troubles precluded a single coherent air campaign in Vietnam, there were promising signs that reflected the benefits of centralization, particularly when it came to the leveraging of theater-wide intelligence products to improve the effectiveness of air targeting. This quickly became evident during the bombing campaign known as Linebacker. “Linebacker, planners worked with a list of approved, validated targets. If a petroleum storage facility was blanketed by clouds, then an alternate target could be struck

¹⁶ The Services had responsibility for combat operations until the passage of the Goldwater Nichols Act in 1964, which established regional commanders with operational control of conflicts in their area of responsibility, and limited the role of the Services to organizing, equipping, and training their respective forces.

(Tilford 236).” This compared favorably with experiences in a prior operation, Rolling Thunder, where poor weather over a target site would result in a ‘scrub’ of the mission – a situation wherein the aircraft returns to base with its weapons undelivered.

This improvement in operational capability is credited to the establishment of a weapons control center code-named Teaball (Tilford). Teaball was established because those at the unit level who would need to task an aircraft with up-to-date targets were normally not cleared to use the intelligence products from certain intelligence collection platforms. These platforms were designed to feed an organization that demanded vertical information flow, allowing the commanders and headquarters staff to see the battlefield situation, and make strategic determinations regarding the use of forces. Teaball was the first U.S. effort to use this information at the warfighter level. “Using a combination of radar and other intelligence-gathering sources and the down-linking capabilities on aircraft platforms like the KC-135 Combat Apple and Olympic Torch or the EC-121 code-named Disco and the Navy radar picket ship called Red Crown (all involved in gathering or monitoring various electronic signals), up-to-the-moment information was sent to a central clearing house...There Teaball analysts used all the information to plot and track enemy aircraft so that tactical decisions could be made based on the latest intelligence. Then, Teaball controllers passed the information that was needed directly to the aircrews (Tilford 242, 243).”

Nevertheless, Carpenter finds that hostility to central air planning persisted long after the Vietnam Conflict, as demonstrated in this quote from the "Operation DESERT SHIELD/DESERT STORM After Action Report" by the Marine Liaison, CENTAF:

“During Desert Shield/Storm it was apparent that the Marine Aircraft Wing was reluctant to become part of the overall air campaign in concert with the other theater air assets. Much of this was due to the inherent fear of the Air Force control fostered by Southeast Asia, and the need to demonstrate MAGTF¹⁷ control over its own air assets.”

Operation Desert Storm and Instant Thunder

Winnefeld, et al., provide a detailed inter-Service history of air doctrine in their history of air power in the Gulf War. While the United States arguably fields four air forces, these

¹⁷ Marine Air Ground Task Force – Organizational unit for Marine fixed and rotary wing aviation

authors indicate that “alone among the services, the Air Force had a theory and doctrine of air power. Moreover, in the Air Force’s view this doctrine applied to all air forces regardless of service or nationality (18).” The Navy had a naval doctrine, with no separate air doctrine, and the Army and Marine view of air power is little changed from the roles of reconnaissance and close air support. The Air Force, in stark contrast, had detailed writings concerning how the parallel, precision bombing of strategic and operational targets could induce a “paralysis” on the enemy forces. As the nation moved towards conflict in 1990, “it was becoming clearer that an era of centrally controlled or coordinated joint air operations was closer at hand than at any time since World War II (Winnefeld, et al., 19).” Air Force Colonel John Warden is credited with providing the framework for the Gulf War air campaign in his timely work in 1989. His theory and the pursuant doctrine called for attacking the leadership of an enemy through the careful selection of targets within several concentric rings. “Command can be attacked in the following three spheres: the information sphere, the decision sphere; and the communications sphere (Warden 45-46).”

In what is still considered a somewhat controversial move (the Air Staff is not usually involved in operational details during a campaign), the Air Force encouraged Col. Warden, at that time head of the Air Force’s “Checkmate” think-tank operation with the Air Staff, to develop a target list that reflected a strategic campaign, based on the notion that air power alone may expel Iraqi forces from the Kuwaiti Theater of Operations (KTO). This targeting plan, dubbed “Instant Thunder,” eventually became the basis for Gen. Schwartzkopf’s (Commander-in-Chief, Central Command) air campaign within the KTO. “In some respects we had learned many of the substantive lessons of Vietnam, although at the beginning of August 1990 Tactical Air Command proved incapable of conceiving of any air role for its forces other than serving as the Army commander’s long-range artillery. Much of the credit for the actual success of effort is due to obscure battles fought in the Pentagon in 1985 and 1986 for a joint air-component commander (Murray 109).”

The prosecution of the air campaign in Operation Desert Storm featured centralized planning and control. The Joint Force Air Component Commander (JFACC) was U.S. Air Force Gen. Horner. While the JFACC role can be filled by a flag officer (one- through four-star general or admiral) from any Service, the intent is for the JFACC to represent the

preponderance of air assets engaged in the conflict. In a largely naval conflict or during the early phase of a conflict where the bulk of the air campaign is executed by Navy aircraft, the JFACC may well be a Navy Admiral. Since 1991, The Navy has developed concepts regarding the “JFACC Afloat,” an air command center aboard the command ship for a carrier battle group. However, for Operation Desert Storm, the bulk of the air assets were U.S. Air Force and the Air Force therefore staffed the JFACC billet. Although the Service affiliation of the JFACC seems a small matter, this individual develops the campaign, strategy, and rules of engagement (ROE) for the prosecution of air warfare. This centralization of the planning and execution function can lead to inter-Service conflict.

In the case of Desert Storm, for example, the U.S. Navy wanted to extend their usual “beyond-visual-range” rules of engagement – which allowed pilots to identify a target and engage from outside their visual range – over the Kuwaiti theatre of operations. Gen. Horner mandated that the ROE (rules concerning when and how U.S. and Allied pilots will fire upon the enemy) included a restriction on “beyond visual range” contacts. The pilots were ordered to have two independent means of verifying the target as an enemy aircraft. Because the U.S. was using stealth aircraft, the fear was that these pilots were at risk if the ROE allowed for an engagement following instrument data alone. Further, he felt the Iraqi Air Force did not pose a sufficient threat that pilots would need the additional latitude provided by the Navy approach. Air Force F-15’s were equipped to operate under this constraint, but Navy F-14’s and F-18’s were not. Some Navy aircraft were equipped with Phoenix missiles capable of engaging an enemy target from 55 or 60 miles away. Also, the Navy pilots trained and performed under a much less restrictive ROE – since a single enemy aircraft can do much damage to a carrier, these pilots are cleared to engage from a much longer range, not allowing the enemy to get too close to the battle group. As a result of this restrictive ROE, the Navy aircraft felt they were overly constrained (Carpenter).

“Our rule was you had to have two separate, independent, physics-based ways of identifying the guy as hostile before you could shoot him. The problem is, the F-14 and the F-18 have only one way you can do it. Of course, the F-15 has several ways you can do it. My perception was the Navy thought the reason we were insisting on two independent means of verification was because we were going to take this opportunity to wrest the Top Gun medal away from these guys. It was a manhood thing.” - Major General John Corder, Deputy Commander for Operations, CENTAF [Central Command Air Forces] (Carpenter”).

So strong was this perception on the part of the Navy that senior officers convinced Adm. Stan Arthur to raise the issue with Gen. Horner, and to request a more “relaxed” ROE in order to allow Navy assets to more fully participate. The issue was brought before the Joint Forces Commander (JFC), who listened to both sides, and decided the more restrictive ROEs should remain, agreeing with the JFACC’s argument regarding risk to stealth and other friendly forces. Nevertheless, Gen. Horner agreed to selectively allow relaxed ROEs so that Navy aircraft were not unnecessarily excluded from pursuing enemy aircraft. This issue became particularly heated after the first night of warfare, when an A-6 Navy plane was lost. Initial reports indicated that the incident was due to friendly fire, and that an F-14 had shot it down – possibly using a Phoenix missile under the relaxed ROEs. The debate continued years after the conflict, with Gen. Corder maintaining that the “navy will never admit it,” and ADM Arthur quoted as saying he “never will forgive” the person who characterized the incident as friendly fire (Carpenter). There were other instances where Navy officers felt various ROEs and procedures were geared to showcase Air Force capability and not the contribution of Navy airpower.

“This war was utilized by the USAF to prove ‘USAF Air Power’, not to prove that combined forces or even joint forces could force multiply and more effectively conduct the war. For example, the F-14 was originally restricted from forward combat air patrol (CAP) positions because CAP aircraft were required to have the ability to electronically identify (EID) and interrogate IFF, friendly and foe.” – Navy representatives to the JFACC Special Planning Cell, as quoted in Carpenter.

The centralization of command here served in this instance to erode the inter-Service trust on precisely the night when the Joint forces needed to coordinate and fight as a single unit. In the case of the U.S. Marine Corps, the trust was not necessarily eroded, because Joint

doctrine was never enacted. The successful personal relationship between the senior Marine (Gen. Walt Boomer) and the JFACC enabled the Marine aviation elements to ignore the ATO planning process altogether.

Historically, since the WWII battle of Guadalcanal when Naval aviation was withdrawn from ground support, the Marine Corps has maintained control of its aviation assets. Marine air is intended primarily to protect Marine ground troops. The aircraft are short range, and their weapons and pilot training are all geared towards this mission. In the JFACC concept, whatever air assets are “excess” were to be made available for use at the discretion of the JFACC. In addition, for Desert Storm, the host nation (Saudi Arabia) required that all Allied aircraft be represented in the ATO. Theoretically, this included the Marine flights that were flown in direct support of the Marine ground troops. The tension was established: the Marine aviation personnel believed their purpose was primarily to support Marines on the ground (the definition of “excess” air assets remained vague), and that they had nothing to gain in submitting targets to the Joint planning cell for addition to the ATO. The centralized air planning function, on the other hand, mandated that the Marine targets be part of the planning process, and that excess Marine aviation be made available to the rest of the ATO targets. Years after Vietnam, where there were lists of “Army,” “Navy” or “Air Force” targets – the evolution of Joint doctrine meant that there were no Service-specific targets or air assets. The following quote from Carpenter demonstrates what happened when this tension was played out in Operation Desert Storm:

“During the execution of the air campaign, it came to my attention that the way the Marines were operating was outside the system that we had established for planning, processing, and then putting information into an ATO before execution. I was bothered by this; for all intents and purposes the Marines were subverting the established planning process.

“I went to Horner and explained the situation. The Marines were bypassing the planning cells where we constructed the Master Attack Plan, which designated targets and force packages to attack them. This occurred about the first week in February. Instead of coming to us (we had a Marine Corps representative in the Iraqi strategic target planning cell) and giving us their inputs, the Marines would withhold information from us. They would go to the ATO cell late at night and give the ‘changes’ to the ATO operators. They would present what they were trying to do as “changes” to the process and give them to the guys processing the ATO.

“So they would accept this information from the Marines as if it were a change, and input it into the system. In fact, it wasn't really a change. It was their initial input. They had to get it into the ATO because they needed the deconfliction, they needed the call signs, the air space management, and so on and so forth. They would bypass the planning cell and go hit whatever they wanted to hit. They could bypass the agreement that was made early on.” – LtCol Dave Deptula, USAF.

Gen. Buster Glosson, USAF, Director of Planning for CENTAF, agreed with Col. Deptula's assessment and frustration regarding the behavior of the Marine Corps in this Joint air campaign:

“The Marines were obsessed with the MEF label. They kept of two thirds of their air assets to support a ground action that was not about to happen and wasn't even in the realm of the possible. They only used one third of their assets to fly sorties that should have been fraggged. This is something that a CINC should have controlled. He should have just slam-dunked them (Carpenter).”

Huber, et al., found that Operation Desert Storm revealed “key shortcomings in current technical capabilities. On more than a few occasions, this conflict served to illustrate the futility of gathering overwhelming quantities of data that could not be transformed into actionable intelligence, delivered to combat commanders where and when needed, and packaged in a form that allowed for immediate exploitation...the needs for information processing far outstripped the capabilities of the infrastructure to meet the demands placed on it (20).” The authors here imply that increased focus on information processing and analysis is needed to enable the planning staff to make use of the information. The partial delegation of planning or target assignment is not envisioned in this work.

Operation Enduring Freedom

The most recent use of airpower in the pursuit of national objectives is the campaign that followed the September 11th attacks on the World Trade Center and the Pentagon by operatives from the al-Qaeda organization. The emerging doctrine from the Bush Administration was to pursue not only terrorist organizations, but “those who harbor them.” After the refusal of the ruling Taliban in Afghanistan to cooperate in apprehending the al-Qaeda leader, Osama bin Laden, the U.S. began combat operations aimed at removing the Taliban from power in order to more effectively destroy the al-Qaeda

organization. Because the Taliban was already facing a rebel insurgency in the Northern Alliance organization, the U.S. coordinated air attacks with the ground incursions to help the Northern Alliance gain ground quickly.

By September 19th, the U.S. had begun deploying more than 100 combat aircraft and several naval carrier battle groups to the region. On October 7th, air operations commenced over Afghanistan. Working with British forces in the area, the initial wave consisted of Tomahawk and other cruise missiles launched from submarines in the Arabian Sea and Persian Gulf. These were followed by carrier-based aircraft and bombers launched from Diego Garcia (U.S. base in the Indian Ocean) and Whiteman Air Force base in the U.S. The initial focus was on degrading enemy air defenses, limited though they were in Afghanistan, and infrastructure targets in the capital city of Kabul and the Taliban stronghold of Kandahar (targets include the Royal palace, television broadcast tower, Afghan Radio, airport facilities, and government offices). Other early targets included al-Qaeda facilities near Jalalabad, as well as a Taliban depot in Mazar-i-Sharif. Military and infrastructure attacks in Herat and Kunduz were also struck on this first night. The country lacked an integrated air defense or competent air force, and air operations quickly became focused on close air support to advancing Northern Alliance forces as well as leadership targets on both Taliban and al-Qaeda targets.

Before the first day's attacks were over, carrier-based fighters were able to engage at a low level on Taliban targets, taking over from the B-1s and B-2s. By October 10th, the U.S. declared air superiority over Afghanistan. By October 16th, AC-130 gunships – aircraft that fly low and slow and are among the most vulnerable in the inventory – were employed over troop concentrations. By the week of October 20th, the U.S. was concentrating on Taliban front line forces in direct support of Northern Alliance troops.

After the fall of Mazar-i-Sharif, there was a pause in the bombing while Northern Alliance forces moved south and west towards southern objectives. As rebel forces began to operate freely throughout most of the country, air operations became focused on suspected leadership positions. Targeting these objectives was facilitated by the presence of U.S. Special Forces on the ground, working with local forces to identify suspect caves and buildings, often using laser designation of targets to guide bombs to these targets.

These forces, proximate to the target, provided intelligence regarding target orientation and defenses beyond that available in most previous air operations. In addition, the RQ-1 Predator and Global Hawk unmanned reconnaissance aircraft provided real-time intelligence to the ground forces. In a new development, Predator video has been “streamed” to the AC-130 gunships, allowing real-time intelligence to the cockpit. In addition, a “new data system enables B-52 bombers to receive targeting data en route to Afghanistan from ... Diego Garcia... allowing the lumbering bombers to fly close-air support missions against ‘emerging targets’ previously flown by A-10 assault aircraft (Loeb).”

Early information implies that bombers were able to enter the area of responsibility, controlled and directed to lucrative targets by ground-based forward air controllers, located within 300 meters of the impact point. The success of the Afghanistan air campaign may appear to prove the U.S. air operations have already broken free of the rigidity of the ATO, and have developed a concept of operations that provides for a self-organized air campaign – where targets are not scripted for attack by a specific aircraft. However, the target-poor environment in Afghanistan (when compared to larger conflicts against adversaries with developed defenses and significant infrastructure targets, as was the case in Operation Desert Storm) meant that the buildup of U.S. weaponry in the theatre quickly became a sufficiently decisive force. Compared to the months of buildup in the 1990 campaign in Southwest Asia, the U.S. was able to assemble an “overwhelming force” against the Taliban within weeks. The favorable ratio of weapons to targets was reached more quickly, primarily because the number of targets was much lower. When the ratio of weapons to targets is not as favorable, air planners are forced to optimize resources in accordance with JFACC phasing. This optimization problem is accomplished by central planning, and reflected in the Master Air Attack Plan. For Afghanistan, planners did not have to solve an optimization of scarce resources.

This development proved useful for Navy carriers, who historically had been unable to receive and digest the ATO in a timely manner. For Operation Desert Storm, the document had to be printed, handed to an action officer, and flown by helicopter to the carriers due to the narrow communications bandwidth available aboard the carrier. Kropf provides this

background from the Korea conflict: “Part of the problem in integrating naval air into the theater air battle was the large amount of communications required by the large, centralized FEAF [Far East Air Forces] system. Carriers had limited communications capabilities, often operated under radio silence, and were unable to handle high-volume FEAF communications. One example of the incompatibility of the high-volume Air Force communications with the limited Navy capacity was a FEAF radio message in November 1950 that gave the air plan for one day. Sent to the carrier task force, it required over 30 man-hours to process (1990).” By 2001, however, the issue seems to have been resolved. The results of reconnaissance missions have been provided directly to the aircraft overhead, resulting in a more timely execution cycle, and in some sense, realizing the promise of this research. However, the lessons of Afghanistan may have less applicability in a conflict on the scale of an Operation Desert Storm or Korea.

Many of the lessons learned in Operation Desert Storm (ODS) are tempered by the observation that an established infrastructure, substantial POL (petroleum, oil, lubricants) stores, and a large coalition force differentiated the ODS experience from future air campaigns. Likewise, the lessons from the Afghanistan campaign must likewise be tempered by the fact that a sparse and largely undefended target set, combined with friendly forces throughout the country, provided an environment that should be considered hospitable to the successful prosecution of an air campaign. It would be premature to assume that the U.S. would gain a rapid victory over air defenses, large friendly forces in country, or freedom of movement on the ground for U.S. Special Forces to aid in air targeting. For example, speculation regarding another campaign in Iraq is accompanied by caution that the air campaign will more closely resemble the 1991 conflict than the ongoing effort in Afghanistan. While some emerging concepts, notably the integration of Predator video feeds with combat aircraft, will doubtless be a factor in the next air war, these assets would encounter more resistance in future encounters than when they were employed against the Taliban forces in 2001.

The Centralized Air Campaign

Our hypothesis suggests that a partially decentralized tasking and execution of air assets may be feasible for the employment of U.S. airpower. This represents a departure from Air

Force doctrine, which has sought to centralize the command and control of air assets since the Vietnam conflict. Excessive decentralization in Vietnam often meant the lack of coordination among Services or areas of responsibility (AORs) within the theater of operations, and Air Force leadership, since that conflict worked to establish the centralization of target selection, tasking, resource allocation, etc. The culmination of their efforts was the establishment of the Joint Force Air Component Commander (JFACC), whom is tasked in a conflict with the design and prosecution of the air campaign, in support of the Joint Force Commander's objectives. In Operation Desert Storm, the initial Master Air Attack Plan, designed by the planning cell supporting the JFACC, was lauded for its tight script and orchestration of the air war; providing for the assignment of specific aircraft to conduct missions along specific routes to destroy/damage specific targets. This approach minimized fratricide from both air and ground assets, assured sufficient refueling, and assured positioning of assets to suppress enemy air defenses (Title X report).

The pilot who found him or herself unable to execute according to the ATO was forced to contact the central planning cell for further guidance. In a bad weather situation, this can lead to hundreds of attack and support aircraft in search of real-time guidance regarding a new target. The attack "package" consists of the pilot, the aircraft and its armaments, its range based on current fuel stores and theater refueling capability, its escort assets (electronic warfare, air-to-air defensive aircraft), and its "time on station" – how long it has been in flight. This package is competing for a new target against other attack packages with different values for these variables. The targets are dynamically entering the master list as intelligence reports arrive regarding the effectiveness of previous attacks, presence of mobile targets, and previously undiscovered enemy assets nominated to the target list.

During the later phases of the Gulf War, two problems with this approach emerged:

1) combat assessments were slow to arrive – leaving the question of re-attack open by the time fixed targets were due to be reassigned; and 2) an increasing proportion of the target list were mobile targets; these are targets whose very nature makes it difficult to plan attack assets with much lead time.

As a result, an increasing number of sorties occurred that were not accounted for in the MAAP. For example, defense suppression flights included not only direct support

operations – escort and suppression in support of a specific strike; and area suppression operations – where several proximate strikes were supported, but also autonomous operations. These last were attacks within specific geographic areas and consisted of aircraft packages roaming kill boxes and attacking targets of opportunity (Winnefeld, et al. 171-172). These sorties, not reflected in the script assigning specific aircraft to specific targets, represent an adaptive response to the inadequacy of the MAAP script to the changing situation.

We can summarize the adaptive behavior of the ODS air campaign planning staff thus:

1. Kill boxes for ground attack fighters
2. Set-aside sorties
3. Aircraft on strip-alert

While this adaptive behavior served the Joint Force Commander well, there may be other approaches to this problem for future operational use. Such an approach is the basis for this research: an attempt to structure a decentralized planning and execution of target assignment to aircraft in a Joint air campaign. Based on the experiences in Vietnam, which led to the development of a centralized planning for air doctrine, the Joint Forces Air Component Commander, we anticipate some cultural opposition to the suggestion of decentralized “planning.” However, if the evolution of command and control communications (C3) systems is intended to deliver increasing information to the warfighter, in order to increase the pilots’ situational awareness (SA), then perhaps the organizational structures may likewise be evolved to leverage this new level of SA. This hypothesis does not suggest the decentralization of air campaign planning in support of JFC (Joint Force Commander) objectives, rather we posit the decentralization of tactical details for the execution of the air campaign plan, specifically details, such as determining which aircraft will strike which targets.

There are two elements to an air campaign worth describing here in brief: allocation and assignment. The allocation of aircraft to specific objectives occurs at the mission type level, and includes the phasing of target types to be struck over time. This is in part done to reflect the JFC objectives, and in part done based on available assets in theater. For example, long range bombers may be available early in a conflict, followed by a small

amount of tactical strike aircraft from naval air carriers, and finally by a larger fleet of short-range tactical aircraft, as in-theater air assets are deployed from the US or nearby permanent bases. The assignment of aircraft to specific targets is done within the Master Air Attack Plan (MAAP) and Air Tasking Order (ATO), and reflects a tactical planning level that may be feasibly bypassed by the use of a self-organizing air campaign – where cooperating aircraft with a common rule-set and picture of the battlefield may “agree” on target/aircraft assignment as the situation develops.

Winnefeld and Johnson make a similar argument while presenting a compelling organizational history concerning the evolution of the Joint Force Air Component Commander (JFACC). While unity of effort generally refers to the coordination of political and military strategies in a conflict, Winnefeld and Johnson restructure the term to include both a unity of command and a unity of control (4). In reviewing lessons learned from their case studies spanning air combat operations from 1941 through 1991, they conclude:

“The CINC [Command in Chief] and his staff should keep out of the details of air operational planning and execution once combat operations have started. But apportionment of effort and approval of target classes are not details. The CINC staff’s job is to understand events and keep the CINC informed, and to insure the air commander receives the proper support (147).” While a unity of command may be necessary to accomplish the intent of unity of effort, the authors argue that control does not necessarily need to be centralized to accomplish this goal (147).

While Air Force doctrine seeks the unity of command and control, this hypothesis is aligned with Winnefeld and Johnson in arguing that a unity of control is not essential to the objective of a coordinated, coherent air campaign – unity of command.

In a RAND report in 1996, Huber, et al., found that information technologies may lead to organizational change for the air component commander. Specifically, they allude to Arquilla and Ronfeldt’s work which observed that future U.S. adversaries were most likely to have a network organization, and that hierarchical organizations encountered substantial obstacles when attempting to combat networked organizations. Building on this, they posit an air staff (supporting functions for the air component commander / JFACC) that is a

networked organization within the hierarchical U.S. combat force entity, resulting in a hybrid organization that realizes the benefits of networks (rapid communication among nodes) while retaining the control necessary within a hierarchy.

Why is central control necessary? Huber, et al., discuss the history of military organizations; drawing a continuum from medieval “personal” combat, where a commander could observe the battlefield and direct resources, through the 17th century, which saw a drastic growth in the area of the battlefield. This growth was due to the increased range and lethality of the weaponry, and was accompanied by communication technology, as well as the development of personal timepieces with which to coordinate attacks (Van Creveld’s observation noted by Huber, et al.). While success in early combat was credited to strength of arms, this growth in the scope of operations meant that success was often a result of better choreography. The victor was often the one who could organize his campaign, outflank the enemy, synchronize attacks to overwhelm, and exploit emerging weaknesses. Thus the commander sought to maintain his personal “presence” on the battlefield, observing the situation and coordinating resources through the use of information technologies. “This opening of the battlefield...has decreased a commander’s ability to oversee the battle and provide direct orders. While the military’s offensive capability has increased, the commander’s sphere of influence has failed to keep pace. He therefore has grown more dependent on assistants to help manage and coordinate disparate activities (Huber, et al., 14).”

Much of the design for the implementation of information technologies in U.S. military operations comes from legislation and Joint Staff initiatives to leverage the “information revolution.” Admiral Owens, then Vice Chairman for the Joint Chiefs of Staff, is quoted in Huber, et al., outlining the role of information in conflict:

“It is the system that converts the information derived from battlespace awareness into deeper knowledge and understanding of the battle space and involves everything from automated target recognition to an understanding of the opponent’s operational scheme and the networks he relies on to pursue that scheme. It is also the realm in which target identification, assignment, and allocation take place. In sum, it converts the understanding of the battlespace into missions and assignments designed to alter, control, and dominate that space (Owens, as quoted in Huber, et al., 19).”

General Colin Powell, former Chairman, Joint Chiefs of Staff, is quoted more succinctly in Huber outlining the goal for advanced communication and information technology on the battlefield: “The ultimate goal is simple: give the battlefield commander access to all the information needed to win the war (21).” Powell’s goal here reinforces the continued centralization of planning with the battle commander.

The Air Force Scientific Advisory Board reflected the Joint Staff position in a list of objectives termed “C4I for the Warfighter.” (C4I refers to command, control, communications, computers and intelligence; and is an oft-used acronym to refer to the uses for information on the battlefield. Variants include C2 – command and control, and C3 – command, control and communications.) In this list, the Board refers to “horizontal/vertical C2,” implying that information could be integrated “up, down, or laterally through the war fighting force or other organizations.” While the lateral nature of information, moving across nodes, is touched upon, the actions taken by nodes based on that information is limited based on the hierarchical organization that is enforced by battle management (Huber, et al., 22).

JFACC and DARPA

The JFACC is the entity charged with applying air assets to achieve the Joint Force Commander’s objectives. The processes are evolving, with a current emphasis on the development of planning tools that incorporate the flow of information from advanced sensor and communications technology into a semi-automated planning process. The main difference between the JFACC approach and the one in this paper is the JFACC’s continued emphasis on the centralization of target/aircraft assignment. Where we posit that increased information can lead to a decentralization of target assignment, the Defense Advanced Research Projects Agency (DARPA) investigations use increased information within planning tools to enable the JFACC planning staff to gain a better understanding of the battlefield.

DARPA is investigating several decision support tools and processes to enable the development of effective ATOs in response to JFACC guidance. These are expressed in a collaborative concept called the “Joint Air Campaign Tool.” DARPA envisions using these tools to provide combat assessment, logistics support, battle management, and aircraft

and weapons analysis (for target assignment). The term collaboration is used to indicate the sharing of incoming information among battle staff personnel and planning tools; sharing that is embodied in shared simulation tools, used for ‘what-if’ analyses of various courses of action. This emphasis on high-speed analytic tools is a top-down approach at managing the dynamic battlefield in an air campaign, while our approach is bottom-up – enabling the individual aircraft or strike packages to select and engage targets based on a central allocation of effort and target prioritization.

Agent-Based Modeling

Agent: an entity capable of election, or choice. Thomas Aquinas

The field of agent-based modeling is relatively new, and represents an attempt to understand certain systems from the “bottom-up,” that is, by modeling components and representing their rules for interaction – then observing the ensuing aggregated behavior. These systems may occur in nature, and include immune systems, multicellular organisms, ecosystems, and other systems that include the interaction of sentient or otherwise individually motivated agents. The aim is to explore system characteristics that emerge as the result of “nonlinear spatio-temporal interactions among a large number of components or subsystems (Tessier).” The mechanism for simulating these interactions is to model individual agents in a common environment, motivate the agents according to objectives and constraints, and observe the interactions – both among agents and between agents and their environment.

We refer to an “emergent” air campaign, referring to one that is made up of *post hoc* observations of the sum of agent interactions. “Emergence occurs when interactions among objects at one level give rise to different types of objects at another level (Gilbert & Troitzsch: 10).” Holland refers to emergence as “complex large-scale behaviors from the aggregate interactions of less complex agents (1995:11),” and the “most enigmatic aspect of [complex adaptive systems]. 1995:12.” This definition of the air campaign does not diverge from history, since the historical treatment of the air campaign focuses on what happened, not necessarily what was planned.

Increasingly, this modeling approach is being applied to the abstraction of large systems, particularly when modeling the interaction of disparate systems. The premise here is that by imbuing individually motivated agents with certain rules and objectives, the observation of agent interactions leads to insights regarding a system that may not be known *a priori*. In fact, the selection of this approach constitutes an assumption that the investigator lacks sufficient understanding about the problem to represent key variables and relationships in a closed equation.

For modeling combat operations, the traditional model paradigm is one of equation-based models (EBM), and therefore a direct comparison of agent-based with these is in order. For equation-based models, the unit of analysis is often groups of individuals. EBMs are often deterministic, with the introduction of random values for selected values over a series of iterations used in order to approach a stochastic analysis. EBMs feature a known environment, represented by established variables, limited domains for variable values, and assumed relationships among these variables. For agent-based models (ABM), the units of analysis are individual adaptive agents. By nature, the analysis is stochastic, as relationships are observed *post priori*. The environment is therefore not completely known, and domains for variables are discovered through agent interactions. The assumption is that by modeling individually-motivated agents as they interact with other agents and their environment, we discover the nature of the system in question. EBM assumes the system and captures these assumptions in a series of fixed equations. Values are randomized and the resulting landscape represents the system's various permutations. For ABM, the resulting landscape represents agent transactions and movements through their environment. For authors in this field, the comparison between EBM and ABM is key: "We consider this feature, the representation of organisms by programs, to be the defining feature of 'artificial life' models of population behavior, the property that distinguishes them from other mathematical or computational models of populations (Taylor & Jefferson, 5)."

Holland asserts that the systems we attempt to understand through the use of models are best thought of as complex adaptive systems, which "abound in nonlinearities...Nonlinearities mean that our most useful tools for generalizing

observations into theory – trend analysis, determination of equilibria, sample means, and so on – are badly blunted (1995: 5).” While equation-based models assume a known relationship among variables, and explore correlations, agent-based simulations model processes, and feature emergent behavior, described in most of the relevant literature as behavior that is described by characteristics not found in the component elements. Multi-agent systems are denoted by a formal description of the agent rules and of the environment. Unlike linear models, a single math formula, however complex, fails to capture the interactions among the agents.

The discipline of economics has been an early adopter for agent-based modeling. Axtell (2000) details two major fields of economic modeling: that which accepts the rational agent assumption and focuses on “bringing new optimization techniques to classic economic models,” and agent-based models. According to Axtell, an “agent-based model consists of individual agents, commonly implemented in software as objects. Agent objects have states and rules of behavior. Running such a model simply amounts to instantiating an agent population, letting the agents interact, and monitoring what happens. That is, executing the model – spinning it forward in time – is all that is necessary in order to ‘solve’ it. Furthermore, when a particular instantiation of an agent-based model, call it A , produces results R , one has established a sufficiency theorem, that is, the formal statement R if A (2).” This last observation refers to a 1972 work by Newell and Simon, in which the authors note that sufficiency is the first requirement of a theory.

Uses for Agent-Based Modeling

Agent-based models have been used to explore issues of decentralized control (can a global strategy be delegated to agents whose individual strategies then serve the global interest) and emergent behavior (what new global behaviors may be observed given the interaction of cooperative intelligent agents). The focus for this paper is the former, where we posit that the “global strategy” of an air campaign can feasibly be pursued through the interaction of cooperative intelligent agents with appropriate individual strategies. Axtell posits a slightly different set of uses for agent-based modeling: from use as a presentation aid for an equation-based simulation, to shed light on partial reference models, and for “important classes of problems from which writing down equations is not a useful activity.

In such circumstances, resort to agent-based computational models may be the only way available to explore such processes systematically (2000).”

Gilbert & Troitzsch provide a history of simulation methods, detailing the uses of equation-based modeling and agent-based simulations. The authors show how multi-agent models, the choice for our work, are derived from the fields of cellular automata (CA) and artificial intelligence (AI). Cellular automata consist of grids wherein cells change their attributes based on the behavior of neighboring cells. Cellular automata are arranged in a grid, and feature identical cells of various states. The states for each cell change based on the state of neighbor cells and upon its previous state. These simulations are useful in observing phenomena such as insect swarms, gossip, etc – where the interactions are local (Gilbert & Troitzsch, 121-122). The most well-known implementation of a CA is Conway’s Game of Life, featuring interesting patterns that result from identical cells with extremely simple rules: the cells flip from white to black (and back again) based on the color of their neighbors. By randomly assigning different initial state colors to the cells, repeatable multi-agent patterns are observed. For multi-agent systems, the underlying environment is often implemented through the use of CA methods.

Artificial intelligence is a vast field that generally refers to the ability of a computer program to learn from interactions, changing its model of the perceived environment and its behavior. Maes differentiates between adaptive autonomous agents and “traditional AI” research: “Traditional AI has focused on ‘closed’ systems that have no direct interaction with the problem domain about which they encode knowledge and solve problems. Their connection with the environment is very controlled and indirect through a human operator (137).”

Architectures and Methods for Agent-Based Modeling

Holland provides a framework for agent rules, which define the agents, their performance objectives, and constraints:

1. “The rules must use a single syntax to describe all *cas* [complex adaptive system] agents.
2. The rule syntax must provide for all interactions among agents.

3. There must be an acceptable procedure for adaptively modifying the rules (1995: 43).”

No matter what method is chosen, the use of a model to abstract reality represents a trade-off between excessive details for fidelity and high-level abstraction for breadth. Neither method is a satisfying substitute for scientific observation, but social science is not a field that lends itself to repeated examination of phenomena in “clean” laboratories. Gilbert summarizes it thus: “In general, accuracy (in terms of data points and assumptions built into the model) is important when the aim is prediction, while simplicity is an advantage if the aim is understanding (18).” The aim of this research effort as an existence proof is not prediction; we do not hold that our results represent what will happen in a non-hierarchical air campaign, rather we show that our results represent a reasonable approximation of an air campaign – simplicity in the model design is therefore our approach.

Many agent simulations are extremely simple systems. They have sensors, through which they detect changes to their environment (and pick up messages from other agents, transmitted through the environment), rules, which determine their reaction as the environment changes are matched to performance goals and appropriate activities, and effectors, through which they influence their environment. Often, the effector involves simple movement across the environment landscape. These three aspects describe dynamically coherent agents, ones whose behavior can be predicted based on a simple calculation of the environment and the rule-set. In order to gain this dynamic incoherence, some investigators add a stable memory to the agents to capture their individual histories. This allows the agent to compare changes to the environment to their internal rule-set, but also allows for a calculation given the agent’s history. One example may be an agent who ‘learns’ that a promising target is actually a decoy, adding nothing to the outcome of the campaign. As a target appears in that location subsequent to this event, the agent may choose to bypass that target, ‘learning’ that it is probably a decoy. From the outside, the agent’s decision to bypass the target seems incoherent; only a consideration of the individual history reveals the cause for the behavior. Adding a stable memory to agents, therefore, brings a measure of unpredictability, and perhaps a more faithful representation of complex adaptive systems than can be found through equation-based modeling. Rocha

refers to systems with this dynamic incoherence as “strong sense of agency” systems, and uses “weak sense of agency” to refer to more predictable agent-based systems, such as robotics or state-determined automata. Increasing the level of dynamic incoherency indicates a stronger sense of agency. This is intellectually appealing, as we can say that human behavior is extremely incoherent in that the reasoning that underlies decisions is rarely completely observed.

Gilbert & Troitzsch assert that there is no generally accepted definition of ‘agent’: “the term is usually used to describe self-contained programs which can control their own actions based on their perceptions of their operating environment (158).” Writing from an artificial intelligence background, Wooldridge and Jennings specify what they describe as a “weak notion of agency” – having autonomy, social ability (interaction through some language), reactivity, and pro-activeness – and a “strong” notion of agency – one which essentially maps these abilities onto human instantiations of similar behavior, what they call an intentional stance, consisting of knowledge, belief, intention and obligation (1995: 4-5, 7).

Therefore, the investigator may speak of the agents’ “belief systems,” or “environmental knowledge.” For our purposes, we detail the intentions of the agents to service as many high-value targets as feasible, working for a common good, which equates to an effective, efficient air campaign. This is anthropomorphic, to be sure, but is in line with existing interpretations for agent-based modeling.¹⁸

The approach for using simulation involves a three-step process to gauge the fidelity of the model to the process under study: verification, validation, and accreditation. Verification refers to the process of debugging the code to insure the model is faithfully reproducing the algorithms and data elements intended by the investigator. Validation refers to the comparison between the model’s output and “expected” or observed results. Often historical data is used to see if the model can faithfully reproduce events. This is somewhat of an art form, since the variables that influenced the historical event may not be

¹⁸ Holland incorporates Bohr’s correspondence principle in his theoretical approach to the theory of complex adaptive systems. The use of a correspondence principle “forestalls...errors [that] occur when the mapping between a simulation and the phenomena being investigated is insufficiently constrained, allowing the researcher too much freedom in assigning labels to what are, after all, simply number streams in a computer (1995:171).”

captured sufficiently in the model to provide replication of observable output. Accreditation is a term used often within the Defense Department, and refers to the rare status of a model that has been “certified” as representing real-world phenomena. Lacking a central accreditation, this step is rare, and would be questioned even if the investigator or model developer claimed authentication.

Sample Agent-Based Models

Economics: Trade Network Game Laboratory

In work submitted to IEEE Transactions on Evolutionary Computation, McFadzean, et al, (2000) outline an application using agent-based computational economics (ACE). In “A Computational Laboratory for Evolutionary Trade Networks,” the researchers propose a laboratory for the “investigation of evolutionary trade network formation among strategically interacting buyers, sellers and dealers.” By focusing on labor market experiments, the researchers hope to facilitate access to agent-based modeling techniques, methods, and visualization among social scientists (2001).

Health Care: CoDiT and the SMASH Project (Martín, et al)

The researchers here sought to advance the notion of agent competence, in that individually motivated, cooperating agents were self-aware when comparing their competence to other agents. Because of the agent’s internal model and awareness of neighboring agent capabilities, the agent is better equipped to request the appropriate cooperative activity from the appropriate neighboring agent. The overall aim of this research effort was to promote distributed case-based reasoning applications for use across domains. The application was CoDiT, a multi-agent simulation diabetes therapy application.¹⁹ This system “consists of a group of agents that perform case-based reasoning (CBR) and are able to communicate and cooperate for the purpose of recommending a therapy (Martín, et al. 4).” The problem set is interesting here, as each agent represents a portfolio of case data from specific patients. As the universe of agents lacks the access to all patient data (due to privacy concerns), the agents must cooperate to share information and recommendations without sharing the raw patient data.

¹⁹ Systems of Multiagents for Medical Services in Hospitals (SMASH), more information is available at <http://www.iia.csic.es/Projects/smash/>

Intelligent Manufacturing: National Institute for Standards and Technology (NIST) / Advanced Technology Program (ATP)

The U.S. Commerce Department established the Advanced Technology Program in 1994 to fund (on a cost-sharing basis) high-risk research in specific areas, including the integration of manufacturing applications.²⁰ One thrust here is to link agile planning systems with the transaction data available from manufacturing executions systems (MES). Selected projects making use of agent-based modeling include:

“A Distributed Agent-Based Lookahead Strategy for Intelligent Real-Time Decision-Making in Manufacturing: Develop a software technology based on lookahead strategies -- strategies that follow a tree of possible resultant events -- to permit real-time decision-making in automated manufacturing systems based on data from the shop floor

“An Agent-Based Framework for Integrated Intelligent Planning – Execution: Develop technologies for a plug-and-play framework of integratable business objects and software agents to enable agile manufacturing by making shop-floor status and capacity information available in real-time throughout an enterprise.

“Agent-Enhanced Manufacturing System Initiative: Design and validate a distributed computer infrastructure for defining and deploying software agents to improve the productivity of semiconductor factories by 5 to 10 percent, an impact worth an estimated \$300 million to \$400 million per year per factory.

“ANTS Scheduling and Execution System: Validate a distributed computer system for factory scheduling that will increase the flexibility, responsiveness, and international competitiveness of U.S. shipyards, resulting in annual savings of \$100 million. [ANTS - Agent Network for Task Scheduling]”

Applicability of Agent-Based Models to Social Science Problems

One concept underlying the use of simulation in the social sciences is that complex behavior can be observed among agents who are programmed with very simple rules. The assumption is that a bottom-up abstraction of agents or components of social interaction (groups or individuals) will yield insights regarding coordination, coalition-building, cooperation, and organizational behavior.

The classic observation here comes as a result of the work done by Reynolds, as he labored to construct a computer simulation of a flock of birds. Rather than attempt to

²⁰ A listing of funded projects under the ATP is available from the NIST web site at <http://jazz.nist.gov/atpcf/prjbriefs/listmaker.cfm>. These four examples were obtained through querying this online resource.

capture all the observed properties of a flock in order to enable it to deal with every conceivable situation (an untenable proposition, Waldrop explains, in an analogy regarding efforts at computer program debugging (Waldrop 281)), Reynolds instead programmed the individual birds ("boids"), and instilled three simple rules: maintain a certain minimum distance from neighbors, match the velocity of those around you, and always move towards the perceived center of the flock. With these simple rules, the resulting 'flock' exhibited extremely flock-like behavior. It seems a flock is an emergent property of a group of flying birds, following a few simple rules.

Gilbert & Troitzsch argue for the use of simulation rather than equation-based modeling for exploring issues within social science:

“[C]omputer simulation has a similar role in the social sciences as mathematics does in the physical sciences...There are several reasons why simulation is more appropriate for [formalizing] social science theories than mathematics. First, programming languages are more expressive while also less abstract than most mathematical techniques, at least those accessible to non-specialists. Secondly, programs deal more easily with parallel processes and processes without a well-defined order of actions than systems of mathematical equations. Third, programs are (or can easily be made to be) modular, so that major changes can be made in one part without need to change other parts of the program...Finally, it is easy to build simulation systems which include heterogeneous agents, for example, to simulate people with different perspectives on their social worlds, different stocks of knowledge, different capabilities and so on, while this is usually relatively difficult using mathematics (Gilbert & Troitzsch 5).”

While Gilbert & Troitzsch seek to distance social science simulations from mathematics, Holland offers the framework for a mathematical formalization, recognizing competition and recombination result in nonlinear trajectories that present a challenge to traditional equation-based modeling: “A successful approach combining generating functions, automatic groups, and a revised use of Markov processes should characterize some of the persistent features of the far-from-equilibrium, evolutionary trajectories generated by recombination (1995:170).”

We use agent-based simulation here in an attempt to observe the emergent adaptive behavior that may occur – behavior that cannot be modeled in conventional formulaic models unless the behavior is predicted. Since our interest is in examining the feasibility of

a non-hierarchical design for tasking unassigned targets to surplus capability, we are interested in observing excessive inefficiencies or ineffective targeting. A simple test for our model may be: Are the targets eventually killed, or do the agents enter into fruitless cycles of negotiation until they must “return to base” because of diminishing fuel? Traditional equation-based air engagement models contain multiple assumptions regarding relationships among agents, some of which may affect the observation of the agents’ behavior. These models are appropriate to examining current air tasking doctrine – which features a centralized command and control of forces – but are not helpful in examining the feasibility of a self-organizing, decentralized air campaign.

Air Force Initiatives in Agent-Based Modeling

The Air Force has begun to apply some agent-based modeling research against the specific problem of air tasking. None of these efforts, however, looks to the feasibility of an auction-based system, with capability packages bidding on targets. Rather, the first explores a continuous air tasking construct, which allows for a more effective centralized air campaign; the second is an insect/pheromone analogy to the assignment of targets to aircraft, while the Agent-Based Modeling and Behavioral Representation Project aims to increase the fidelity of controller activities within Joint simulations for wargaming and analytic purposes.

Continuous Air Tasking

Under the Joint Air Operations Functional Process Improvement Project, the Headquarters, Air Combat Command Director of Requirements has identified several functional process improvement areas, among them, the need for “Continuous Air Tasking (e.g., rolling Air Tasking Order (ATO), dynamic ATO).” When combined with the other areas in this study, the improvements were expected to enable “better tracking of target data with air tasking, more reliable terrain data, and more timely updates of air traffic control displays. The result will be more effective use of air combat power.”²¹

²¹ Information taken from Project web site, found at <http://www.c3i.osd.mil/bpr/bprcd/3404.htm>

Insect Pheromones

The Defense Advanced Research Projects Agency (DARPA) in 2000 funded a study that will investigate a self-organizing air campaign. Using an agent-based approach, this study will attempt to replace all planning from the JFACC point of view by placing sensors on the battlefield that release “pheromones” when they detect enemy activity of interest to attacking aircraft. The aircraft, hovering nearby, are alerted and respond to the pheromones. This approach to self-organization uses an insect analogy, doing away with a hierarchical approach to the air operations problem by replacing command hierarchies with a combined visualization.

While a swarm analogy may be technically correct, the execution of an air campaign also involves control mechanisms for responding to contingencies, for example: cessation of hostilities, or changes to operational situation that requires a massive re-tasking of air assets (caused perhaps by an enemy breakthrough). A pure decentralized construct may reduce the probability that such control mechanisms would be immediately effective. The fact that any one of these attacks may have “strategic effects,” such as was the case with the Al Firdos bunker in Operation Desert Storm²² or the bombing of the Chinese Embassy in Yugoslavia in 1998²³, means that control mechanisms are essential to the successful prosecution of the air war.

This interesting research concerns the use of adaptive autonomous agents in air warfare, apparently without the benefit of any human planners. This extreme swarm method of an air campaign will not be presented as a recommendation from this research. Our semi-hierarchical approach to an air campaign is rooted in sound logic and rationale, by retaining the central control of resource apportionment, the key to the strategy of an air campaign. By eliminating the human, allowing each individual agent to sense and destroy its own target, this approach appears to diverge from the ability of the planner to employ phasing over the period of the air war. While our approach allows for a variance in priorities, this

²² The Al Firdos bunker was functioning as a civilian shelter at the time it was struck by allied missiles, leading to loss of civilian life and an opportunity for Iraq to decry the war effort.

²³ An outdated map of Yugoslavia led to the destruction of a building housing the Chinese embassy, and not the desired target.

swarm method is singular in focus and apparently does not allow for any prosecution of a grand strategy other than “kill what you find.”

Agent-Based Modeling and Behavioral Representation (AMBR) Project

The Air Force Research Laboratory, Human Effectiveness Directorate, Sustainment and Logistics Branch, has been investigating the use of agent-based simulation for their mission. In March, 2001, they released a contract solicitation that seeks the development of a “Intelligent Controller Node (ICN)” in order to facilitate the depiction of air controller tasks in the Joint Synthetic Battlespace: “[T]hree types of ICNs will be developed, enabling: (1) the modeling of command and control echelons, (2) the performance of technical controllers or support cell operations, and (3) the simulation of complex human behavior (Commerce Business Daily, March 23, 2001).” This project highlights the continued interest within the Air Force in improving battlefield processes.

The AMBR initiative comes under the Human Effectiveness Directorate, where researchers are working with agent-based simulations to emulate intelligence in simulations. The ABMR project seeks to improve the representation of the air controller process within a battlespace environment, effectively capturing the existing hierarchical, centralized process for presentation within Joint simulations.

3. APPROACH

Research Approach, Design & Method – The ABACUS Model

Research Approach

In order to simulate the observable phenomena that may accompany a self-organizing air campaign, we turn to the modeling techniques specifically designed for the study of complex adaptive systems. Equation-based modeling assumes knowledge of the variables and their interactions, as well as a range of results. Unanticipated results are normally obtained (in computer simulations) using random number generation to obtain random data values – but values still occur with a fixed domain. Using agent-based simulation methods involves not an equation, but the establishment of an environment and autonomous agents – each with objective functions and rules. Maes indicates that the “main problem to be solved in autonomous agent research is to come up with an architecture for an autonomous agent that will result in the agent demonstrating adaptive, robust and effective behavior (138).” Emphasizing the applicability of this approach when examining a decentralized control problem, the author continues: “agent architectures are highly distributed and decentralized. All of the competence modules of an agent operate in parallel. None of the modules is ‘in control’ of other modules...Because of its distributed operation, an agent is typically able to react quickly to changes in the environment or changes in the goals of the system (142).”

The design of this research effort is to employ “agent-based simulation” to simulate the activity of agents bidding on target sets, given near-perfect information regarding their environment. The perspective here is to take a theater-wide view of situation awareness, wherein what is known about fixed/known and mobile/emergent target sets is known fully to each warfighter. We examined the time to clear for such an “auction,” in order to

present a new way of prosecuting the air campaign in Joint operations – one that leverages the use of advanced information technologies in order to realize efficiencies.

Research Designs & Methods

In comparison to the prevalent use of AI or adaptive autonomous agents in modeling, we do not seek in this research to learn more about the “agents”; we leave the detailed examination of pilot motives and adaptive behavior to future research. Much of the social science literature examines how this research method can help learn more about the behavior of the agents whom we abstract in our models. Referring to the emergence that can be observed as a result of interaction dynamics, Maes observes that “ethologists have stressed that an animal’s behavior can only be understood (and only makes sense) in the context of the particular environment it inhabits (140).”

Research Method

The simulation is an abstract of that part of an air war that is necessarily unscripted. When aircraft, assigned to a particular target, are not able for whatever reason to strike their primary (or secondary, if assigned) target, they become unplanned resources located somewhere on the geographic plane. This is represented in ABACUS as the appearance of aircraft on the grid at random times and random places, with random²⁴ distributions of fuel (never less than “bingo fuel²⁵”) and armaments.

The purpose of this paper is to provide an existence proof, and the criteria are therefore simple: Given these rules, will the agents engage and kill targets, the core function of an air campaign? Further, will the resulting emergent air campaign resemble a reasonable use of resources? Future research will address issues of efficiency and possible comparison with hierarchical models of target assignment or historical events. Once the feasibility of this approach is established, it may be that the use of autonomous agents, tasked with decentralized execution within a planning hierarchy, will lead to more flexible and

²⁴ All representations herein of “random” numbers recognize that computers approximate a representation of random behavior, but do not actually provide pure random number generation.

²⁵ “Bingo fuel” refers to the minimum level of fuel required to return to base, as calculated from the aircraft’s present position. Any further movement away from base results in a situation where the aircraft cannot return safely, and alternate landing options or abandonment of the aircraft must be considered.

effective use of airpower in future conflicts where resource optimization is essential to success.

The model used here is ABACUS, written for this purpose. ABACUS is an adaptation of the SugarScape model within the Swarm modeling environment, and provides for agent-based simulations on a cellular automata grid.

Analytical Methods and Models - SWARM

We use the “Swarm” modeling environment (www.swarm.org) developed at the Santa Fe Institute to construct the environment and agents. This simulation system provides a framework within which researchers develop discrete event simulations and has been used in applications in biology, ecology, economics and political science. It has also been used in battlefield simulations (Robert H. Kewley, Jr. and Mark J. Embrechts, “Fuzzy-Genetic Decision Optimization for Positioning of Military Combat Units”, 1998 IEEE World Congress on Computational Intelligence.”). Minar, et al. offer this explanation for the provision of free Swarm code developed at the Santa Fe Institute: “In order for computer modeling to mature there is a need for a standardized set of well-engineered software tools useable on a wide variety of systems. The Swarm project aims to produce such tools through collaboration between scientists and software engineers. Swarm is an efficient, reliable, reusable software apparatus for experimentation (2).”

Daniels identifies two major uses for the Swarm libraries: an empirical evaluation of dynamics, and synthesis. The latter has been applied within basic science applications to extrapolate the “possible” from known biological and chemical interactions. The former is the basis for our investigation, the need for which is further described by Daniels: “[t]he combination of autonomous entities in a shared environment is typically a mutually recursive process that is analytically intractable. In many systems, the only way to know what global dynamics will occur is to run the numbers and find out (1999).”

Building on this conceptual code base, we create agents representing friendly aircraft and enemy targets. The source code for ABACUS is a modification of the SugarScape code developed by Nelson Minar and available online. SugarScape is discussed in chapter 2,

and is a biological model representing agents who “die” without a sufficient regular intake of “sugar” which they obtain from cells on a grid.

Research Design - The Sugarscape Agent-Based Modeling Environment

The SugarScape effort is concerned with answering questions of social organization, examining how social structures and group behaviors arise from the interaction of individuals. The implementation consists of a grid of cells upon a torus, each containing some level of “sugar.” Agents are instantiated upon this surface, with various values for vision, speed, metabolism (consumption of internal sugar), and motivated to obtain sugar in order to sustain “life.” They can “see” the grid around them, to varying levels of distance, and move to the nearest cell that contains sugar. Every time an agent moves, it burns sugar at an amount equal to its metabolic rate, and dies if the internal sugar reaches zero.

Epstein and Axtell motivate their agents by providing rules for their agents: the agents are to move where the environment provides sugar, and “consume” the sugar. “Look around as far as your vision permits, find the spot with the most sugar, go there and eat the sugar (Epstein & Axtell 6).” By then structuring an uneven distribution of sugar in the environment, they authors provide for a dynamic landscape, with agents moving about to maximize their performance function. Add to this mix the concept that without sugar over certain time steps, the agents die, and the simulation begins to take on richness, as characteristics that permit the agents to reach new sugar stores in time (low metabolism, high vision) are preferred in the evolving population.

Functional Description: Sugarscape

In the model, the authors provide for a “landscape,” a two-dimensional grid, where each cell or location has an x and a y coordinate. Each cell contains a current sugar value, a rate of growth (for that sugar value) over time, and a maximum sugar value. As time passes, the sugar value for each cell increases based on the individual growth rate, until the maximum is reached. (One modification of the SugarScape code for our use was the addition of a counter to account for the passage of time. For ABACUS, then, time is indicated by a counter in the code, which increment by one time unit as the landscape cells advance their sugar values according to the initial values.)

Onto this landscape, the authors introduce agents. Agents consume varying units of sugar per turn (the rates differ with each agent, and represent metabolism), and begin with an initial value of internal sugar. When this sugar stock is expended, the agents “die” and are removed from the simulation. The programming logic motivates the agents to survive, and they are thus motivated to travel to nearby cells in order to consume the sugar located there. They can detect the sugar level of cells around them, in a radius determined by their “vision,” also set by an initial value. (This is later used to simulate an evolutionary fact that agents with greater vision have a great chance of survival than those with comparatively limited vision. This observation also applies to agents with lower metabolism, higher initial values, etc.)

The agents can carry unlimited amounts of sugar, and consume all the sugar contained on the cell they currently occupy. (In another modification to this code for our use, ABACUS agents carry a limited amount of sugar for use as they travel, and can collect sugar at a limited rate per turn.) At each turn, the agent looks across its field of vision, and moves towards the cell with the greatest concentration of sugar. Some succeed in reaching their objective and refueling; others run out of time before successfully refueling. Agents are “born” into this world at a certain rate, and either are instantiated with values for speed, metabolism and vision through a randomization algorithm, or by inheriting characteristics from “parents” in simulations that abstract sexual reproduction.

Terna asserts that Sugarscape is a “sum of cellular automata and an agent-based model (2001),” because cells have varying and dynamic levels of sugar. These levels grow over time according to environment rules, and are decremented as agents consume sugar from the cells. The more familiar applications of cellular automata feature an interaction among cells, in that values within a cell are affected by the values in neighboring cells, such as in Conway’s Game of Life. However, if one accepts Terna’s description, Sugarscape represents an interaction between agents and cells. “Technically speaking the agents are instances of object classes; the environment can be interpreted as a cellular automaton, with simple or complicated rules defining the production of food in each cell (2001).”

Beginning with this simple representation of individually motivated agents, Epstein and Axtell go on to build in elements such as spice, which introduces the possibility of trade among heterogeneous agents (with varying needs for sugar or spice). Over the course of their work, they introduce the concepts of conflict, sexual reproduction, tribal relationships, and cultural segregation (Terna).

Modification of Sugarscape for this Work

In the initial design of ABACUS, the initial plan was to use the agents to represent attacking aircraft, and the cells as fixed targets. However, this would have limited the representation of the target set in an air campaign, and rendered the simulation less interesting as a result. Therefore, the cellular automata aspect of Sugarscape is not used in this initial version, reserving it for later use – perhaps for modeling terrain or fixed defenses. Instead, two types of agents, targets and aircraft are allowed. Some targets are fixed, others move with varying rates of speed. Aircraft all move, but with varying rates of speed. These are fixed according to the aircraft type each new agent represents (F-15E, F-16, F-18E/F, or B-2). The landscape is used to provide communication, and for relative geography.

The target set is determined within the configuration text file, where the target classes (surface-to-air missiles, armored personnel carriers, self-propelled howitzers, tanks or C2 node (fixed)) are established by percentage for the battlefield. The attributes for each target class are also configurable, and are speed (number of cells traversed in a turn), minimum and maximum levels of sugar. Sugar is used to indicate the ‘sweetness’ of the target, that is, the priority placed upon it by the Joint Force Air Component Commander. The higher the sugar level, the greater the priority placed upon the target class by air planners. To insure all targets are not identical within a class, the actual sugar values are randomized between the minimum/maximum levels designated in the configuration file. Sugar is not increased or decreased over time; the value is set at agent-target instantiation. For targets that move, they progress in a random direction from their initial position, and generally continue traveling in that direction at their maximum speed, occasionally changing direction also at random. Targets change direction when they reach the edge of the

landscape. Future versions of ABACUS may allow for the simulation of road networks, and confine the moving targets to roads or known paths. Targets do not attempt to evade their attackers, and are not made “aware” of their plight.

Aircraft are instantiated as indicated earlier, according to a certain percentage of each aircraft class. Aircraft carry weapons, which are not displayed in the simulation, but which provide the aircraft with relative capability against the various target classes. The use and effectiveness are tracked and reported, but missile flight is not simulated in the model. The weapons are given “probability of kill (Pk)” percentages for each target class, and each aircraft carries a store of one or more types of weapons. The weapons include the AGM-65 Maverick, Sensor-Fused Weapons within a wind-corrected tactical munitions dispense, Joint Direct Attack Munition, and an Mk-82 gravity bomb. The aircraft loses capability as weapons are expended. In addition to weapons, aircraft have attributes indicating speed, fuel burn rate, and time aloft. The time aloft, combined with the fuel burn rate and time value for each “turn” in the simulation determine how long an aircraft can remain on station before returning to base to land and refuel. Later versions of ABACUS could well feature refueling areas on the landscape, where aircraft could refuel in the air. When the aircraft returns to base, it remains for a given number of turns. This abstracts the real-world variable known as sortie rate – the number of times per day a given aircraft type can generate a takeoff and landing.

The aircraft seek to converge with targets that have the higher sugar values, therefore pursuing targets of higher priority value. This can be configured, so that aircraft begin to prefer targets in proximity over distant “sweeter” targets as their fuel level drops, and they are more pressed to return to base. The landscape cells hold information regarding the target locations, and the current situation regarding targeting by agent-aircraft. This information represents the coordination we have (omit) built into the original Sugarscape code. Sugarscape did not feature communication or coordination among agents, but our (the) simulation required the ability for aircraft to pursue only those targets for which they were better suited than any proximate aircraft. By consulting the information for each target held within the appropriate landscape cell, the aircraft can determine if it is better

equipped to destroy the target than the aircraft currently aiming at that target. If no aircraft has targeted the target, then the landscape cell value is changed to reflect the attacker's information. The aircraft are not **cooperating**, in the sense that they do not bargain for targets, nor do they trade capability; they are **coordinating** according to an agreed set of principles. If a certain aircraft is moving towards a certain target, and another aircraft is a better choice for that target, the first aircraft agrees to seek an alternate target without questioning the situation. Pk values are established in the configuration file for each weapon/target combination, providing for a "killer-victim scoreboard" methodology for adjudicating attacks. When an aircraft chooses to move towards a target that already has an attacker "assigned," three variables are compared to determine who "wins" and gets to continue attacking the target: proximity, Pk, and appropriateness of weapon. The configuration file allows for the analyst to configure the preference among these three comparisons. For example, "Closest, Highest PK, BestFit" in that order would indicate that the aircraft nearest the target wins the contest for the target. In the event the aircraft are equidistant from the target, then the aircraft which has weapons with a higher Pk against that target would win. If Pk values are equal, then the aircraft planning to use the "most appropriate" weapon for the target type wins the contest. In the unlikely event that all these values are equal, then the incumbent aircraft wins, and continues to attack the target.

During each time unit, or turn, each agent engages in the following logic: check remaining fuel to see if they need to return to base, select a target (searching through the entire list of targets in search of the "best" for their situation, coordinating with other aircraft as indicated earlier), and move towards that target at the speed indicated in the configuration file for that aircraft type. When selecting a target, the aircraft checks the target priority (the "sweetest" target is selected), proximity, and then highest Pk (the target they are most likely to destroy is selected). The order in which these attributes are checked differs by aircraft type. An additional layer of complexity can be attained, as each aircraft type can feature a change to the order of these attributes when it is found to be low on fuel.

The use of Sugarscape does not make use of variables or rules that allow for the anthropomorphic representations found in the original work. Rather, the use of agent-

based modeling explores the decentralized execution of a central strategy, and does not need to consider variables that describe human behavior. In fact, it may be that the promise of information and communication technologies combined with organizational constructs that enable self-organization at the task execution level may lead to an increased use of unmanned aircraft for certain tasks.

For these purposes, the agents interact with one another in order to choose among available targets. There is still some centralization to this self-organizing air campaign, however, since each agent relies on a centralized understanding of target attributes. This is properly aligned with the notion that even this self-organized air campaign features a centralized apportionment of assets against certain target classes, which implies a common understanding of target characteristics.

ABACUS

ABACUS is an acronym for Agent-Based Air Campaign Using Swarm. The name indicates the approach to the problem, the problem itself, and includes the modeling language environment.

In its simplest form, an air campaign consists of various types of aircraft, with certain cycle times between launches, weaponry, and fuel capacity. These aircraft pursue and destroy targets, which are prioritized by decision makers not represented in this model. The targets ordinarily have organic and theater defenses, which are also not represented in this model. The model is designed to simulate aircraft that become available after failure to engage their primary or secondary targets. In addition, targets become known to the environment during the simulation, as they are determined to be insufficiently damaged from a previous attack, or emerge from a previously benign landscape.

The environment in ABACUS therefore includes targets and aircraft that appear randomly on the grid. The targets have characteristics that denote their specific vulnerability to various weapons. For the baseline case, target classes include tanks, self-propelled howitzers, surface-to-air missiles (mobile), armored personnel carriers and notional fixed targets (representing above ground buildings such as command and control nodes).

The attack aircraft each have variables that represent fuel capacity and levels over time: weapon types and current stores, sortie rates (the average times per day each aircraft can launch), and percentage of total aircraft over the course of the scenario. The targets themselves are also agents in the model, and their variables include: percentages of total target base (for that class of target), speed (fixed targets are set to zero speed), and a variability of low to high levels of “sugar,” which indicates the relative priority. Finally, the aircraft have certain weapon types, each of which has a different probability of kill (denoted Pk) against target types.

Baseline Configuration File

The following is offered to convey the richness provided by the baseline configuration included in the current version of ABACUS. Most of these parameters can be modified using the configuration file. See Appendix A for an annotated version of the configuration file. The aircraft seek the targets that represent the highest priority for which they are carrying appropriate weaponry. If another aircraft has “decided” to attach that target, the aircraft communicate (through the environment) and compare Pks. The aircraft with the higher Pk continues to attack the target, while the losing aircraft begins searching the environment once more for an appropriate target. If the Pk’s are identical, the aircraft that is closest to the target wins the negotiation and continues approaching the target. This negotiation can occur at any time up until the target is attacked.

At first, the aircraft seek the targets of highest priority, even if their Pks are relatively low against those targets. As their fuel dwindles, and the risk of returning to base with unexpended ordnance rises, their decision logic changes to prefer targets for which they have a high Pk (and therefore a lower probability of being “bumped” by another aircraft. While initially, they may seek a high priority target despite carrying a weapon with a low Pk against that target, the aircraft will change strategy and begin looking for targets they can win in the auction by having a higher Pk. Their overall goal then is to destroy high priority targets, but a secondary objective is to expend their ordnance, and they will pursue lower priority targets as their time runs out. The aircraft return to base when either their fuel level is low enough to warrant refueling, or they have expended their munitions.

Issues

There are several key issues to consider when comparing the ABACUS model to reality. Although the intent here is to demonstrate the feasibility of a non-hierarchical, self-organizing approach to some air tasking (a finding that may be useful in future decisions regarding resource allocation), there are several tactical issues regarding the abstraction of the air campaign that can be addressed briefly.

Pilot discretion and the lack of “perfect” information – While the delivery of information has been automated within ABACUS, potentially to the weapons platform itself, there is ample reason to maintain an ability for the pilot to override the decisions of the system. The decision to fire a weapon at a target, for example, has historically been the province of a human with the ability to reason and assimilate information that the system cannot. While we (one) can assume perfect information within the model, in fact the information available to a decision maker is rarely perfect. Follow-on research with ABACUS may focus on uncertainty, the introduction of error, imperfect or unexpected information, and associated consequences.

There are several examples in recent U.S. military history that bear this out – each a case in which a pilot made a critical decision, and each a case that makes an automated decision to fire somewhat unthinkable. The accidental bombing of the Chinese embassy in Yugoslavia (1998) occurred because the U.S. military used an outdated map that did not indicate the current use of a specific building. The death of Iraqi civilians in the Al Firdos bunker (1991) occurred because a facility associated with military command and control was suddenly used to shelter civilians. While each of these tragedies occurred despite the best efforts to avoid or minimize civilian casualties, the policy backlash was against the use of “fire and forget” or “standoff” weapons – weapons that allow the pilot to engage a target from a distance sufficient as to potentially miss critical clues. A hypothetical example may be the presence of a school bus in a combat zone. Discerning whether the bus was filled with schoolchildren or combatants is a decision (given the limitations of current technology) best left to a human, due to the potentiality for unintended consequences.

Information glut – The delivery of raw data to the cockpit or weapons platform will, in all likelihood, overwhelm the processing capabilities of onboard information systems, even those envisioned in the ABACUS construct. In addition, the processing capability of the pilot, (if the environment is one in which the pilot is permitted to override target selection) is quickly overwhelmed, particularly given the distractions of combat flight and mortal threat.

The availability of increased environment and target information to the cockpit necessitates the integration of automated information analysis to aggregate and filter the information upon which the platform or pilot must base their decisions. As proliferating sensor technologies provide for an increase in information flow, fusion and analytic technologies are required to assist in the assimilation of this information. There are two information processing technologies of interest here arising from the convergence of advances in the fields of artificial intelligence and computational linguistics.

Data Mining. Data mining refers to the use of predictive computer models to discover information from among (usually large) repositories of structured data, that which resides in databases. Examples of automated analytic methods employed to discern or learn patterns in such data include correlation, regression, and link discovery. The use of data mining tools and models would be key to the envisioned information architecture to reduce the flood of information requiring human interpretation.

Text Mining. Text mining, also referred to as text data mining, is a promising field concerned with encoding an understanding of language in unstructured digitized information to allow for the application of data mining techniques on previously “unreachable” information. For example, while data mining models may be applied against a set of data that features a predictable format, with defined rows and attributes such as that found in relational databases or structured spreadsheets; these models cannot simply be “run” against a set of emails, or word processing documents, or news feeds. The information in these resources may be as valuable to the analyst, but finding the individual elements (e.g., names, locations, or organizations) is a manually intensive process. While often referred to as “unstructured” information, in fact there is an underlying linguistic

structure common to these resources: rules regarding grammar, syntax, as well as common usage patterns within a given language. All provide a structure which machine learning specialists and computational linguists can exploit with computer models.

Text mining technologies based on natural language processing read these documents, and extract the concepts based on a learned understanding of linguistic structures. By automatically populating a database with “found” concepts within unstructured information, these technologies provide a basis for data mining technologies, enhanced search capabilities, and a host of information tools with which to tame the information glut that results from advances in information discovery and dissemination. Other text mining technologies rely on a symbolic representation of individual alphanumeric characters, and apply statistical probability models based on the “closeness” of these characters in order to find similar information in large repositories.

The combination of data and text mining provide capabilities for a layer of analysis between the sensor and the decision maker or pilot. This analytic layer represents the added value provided to information repositories by the intelligent use of aggregation, analysis, abstract modeling techniques, and other integrative applications to provide insight to the user. The analytics layer is strongly analogous to the application layer in classic information technology systems, but refers to the applications specifically used to develop insight into the underlying information. Examples of these analytics, employed for various functions, include abstract systems modeling for the exploration of multiple scenarios, discovery of underlying patterns through text or data mining, and link analysis for the discovery of non-obvious relationships among entities within underlying data.

4. RESULTS

Do Unassigned Agents Engage and Kill Targets?




Figure 4-1 Baseline Run – Targets Killed Over Time

TABLE 4-1 BASELINE - WEAPONS EXPENDED

Baseline	609 total weapons				
	AGM-65	TMD-WC-SFW	JDAM	Mk-82	TOTALS
F15E	93	0	0	169	262
F16	29	0	0	82	111
F18EF	0	0	56	62	118
B-2	0	118	0	0	118
TOTALS	122	118	56	313	609

In order to test this thesis, one must test the agent-based simulation's capability to execute a grand strategy in a decentralized command architecture. Rules for agent interaction with

each other and the environment replace scripted orders from higher command regarding agent behavior.

The first of these tests is to explore the model's ability to "kill targets." In this validation step, the results are examined to insure that aircraft pursue targets, select them for destruction, and expend munitions in order to accomplish that goal. Further, the aircraft is expected to "decide" when it is no longer "safe" to remain airborne, and prefer returning to base over further prosecution of the conflict.

In the baseline run, targets are destroyed as aircraft make reasonable use of their weapons, relinquishing their search for the highest priority targets as they "realize" that their fuel stores are low. The result of the ABACUS model is that aircraft, given simple rules based on their fuel and range limitations and the ability of their weapons to destroy given targets, will indeed seek out and destroy targets appropriate to their weapons suite within the time they have left in the "game."

Figure 4-1 shows the Baseline run, with targets destroyed on the y axis, and time represented on the x axis. The x axis is actually referring to "turns," that is, a cycle wherein all agents have considered their position and objective, and have moved towards their objective in accordance with their speed rules for a single turn. In this Weapons Over Time view, we see the accumulation of destroyed targets as the simulation is run. This simple test indicates that targets are being destroyed, as per the overarching objective for the grand strategy that calls for the destruction of enemy targets. In Table 4-1, the weapons are expended by type. For these purposes, the total weapons expended will be used to compare this baseline run against alternate runs. For example, when the effectiveness of the weapons is halved, one expects to see more weapons expended in that scenario as more weapons per target are needed to effectively destroy the target.

Although the baseline run results in an observation that aircraft kill targets according to the priority in the configuration file, additional runs are required to determine the degree to which we find the model a reasonable representation of an air campaign. We select three criteria, posed in the form of questions to the model:

If weapon Pk's are lowered by half, what is the effect on model results? One would expect to see the "war" take many more turns to complete, as weapons do not accomplish the goal of target destruction as they did in the baseline run, and therefore remain on the landscape. Additional resources are then expended to destroy the targets, and a lower y:x ratio is expected when compared to the baseline run.

If fuel capacity on the aircraft is reduced by half, what is the effect on model results? One would expect to see fewer targets destroyed over time, given the reduced range and time on station for the attacking aircraft.

If the target priority is shifted from C2_Nodes to SAMs, what is the effect on model results? One expects to see more SAM targets attacked and destroyed, as aircraft prefer these targets over other classes. This would represent a phase of the conflict wherein the Joint Force Air Component Commander has made the destruction of tactical air defenses his/her priority.

In the sections that follow, the baseline results are compared to each of these scenario changes.

Lower Pk

If weapon Pk's are lowered by half, what is the effect on model results? One would expect to see a reduced effectiveness overall, and increased effort (in terms of multiple attacks) to fully destroy a target.




Figure 4-2 Weapons Effectiveness Halved – Targets Killed Over Time

TABLE 4-2 WEAPONS EFFECTIVENESS HALVED – EXPENDED WEAPONS

Half Pk	1088 total weapons				
	AGM-65	TMD-WC- SFW	JDAM	Mk-82	TOTAL
F15E	137	0	0	260	397
F16	38	0	0	129	167
F18EF	0	0	129	212	341
B-2	0	183	0	0	183
TOTAL	175	183	129	601	1088

With the Pks halved, one would expect to see the effectiveness reduced drastically. Fewer targets would be destroyed over time as the same weapons used in the baseline run were now half as effective – that is, their probability of kill is reduced by fifty percent. Each time an aircraft fires at a target, a random calculation is performed to adjudicate success or failure for the kill. This is standard modeling practice for military simulations; the higher the Pk, the higher the probability that an engagement between an aircraft firing that weapon and a target will result in the destruction of the target. Cutting the Pk in half should result in a longer run (more turns) and more weapons expended in order to kill the

same number of targets. Figure 4-2 shows that indeed a halving of the Pk results in a fewer targets destroyed over the same time period, and Table 4-2 shows the disparity in the total number of weapons expended over the course of the simulation when compared to the baseline run. Whereas the baseline run used 609 weapons to destroy 250 targets, one sees that the Half-Pk run uses 1088, and never kills all the targets within the same time frame as the baseline run (3,285 turns).

Lower Fuel

In ABACUS, the aircraft remain “on station,” or in the air, looking for targets as long as there is sufficient fuel to keep them aloft and return them to base. At each turn, the agents calculate their distance from base, the fuel required to get them home safely, and compares this number to the fuel remaining in the aircraft. When the fuel level approaches this number, the agent changes its objective from killing targets to getting back to base. Of course, each turn also decrements the agent’s fuel stores to simulate the usage of actual fuel in flight. When the agent returns to base, it remains for a set number of turns (based on aircraft type) to simulate the “sortie rate” – the rate at which aircraft can perform missions. For a long-range bomber, the sortie rate may be as low as .5, meaning the aircraft performs one mission every two days. For shorter-range aircraft, this rate may be 2, which would mean that the aircraft departs, flies a mission, recovers, re-arms and refuels, and conducts a second mission all in the space of a single day.

The aircraft begin with varying levels of fuel, and their total capacity varies based on aircraft type. This reflects the “appearance” or more appropriately, the sudden availability of an aircraft onto the landscape as it is no longer flying the scripted route from the ATO. As such, the aircraft may have full or partial fuel levels, but for this simulation, always begin with enough fuel to return to base. The fuel capacity was decremented for each aircraft by fifty percent. This should result in shorter missions for each aircraft. The sortie rates remain the same, so each aircraft will have less time on station than in the baseline run. The result should be a much longer simulation, as many more days will be required to destroy the target set. The number of weapons used should be close to the baseline run, since the engagement parameters (Pk) are unaffected. (The total number of weapons will

be greater, as this campaign will encounter a greater number of targets. Targets are generated throughout the simulation, and a longer simulation will produce therefore a higher number of targets.)




Figure 4-3 Half-Fuel Run – Targets Killed Over Time

TABLE 4-3 HALF FUEL RUN - WEAPONS EXPENDED

Half Fuel	694 total weapons				
	AGM-65	TMD-WC-SFW	JDAM	Mk-82	TOTAL
F15E	86	0	0	165	251
F16	55	0	0	121	176
F18EF	0	0	84	119	203
B-2	0	64	0	0	64
TOTAL	141	64	84	405	694

Prefer SAMs

The aircraft prefer targets with a higher priority (simulated by a higher sugar content using the SugarScape methodology). As such, each run features agents who prefer the command and control (C2) nodes when selecting a target. This is a decomposition of the

grand strategy, which calls for a phased air attack, and early destruction of the enemy's ability to coordinate his forces through the targeting of communications nodes. In order to test the model's ability to reflect a grand strategy, a shift is reflected in that strategy wherein surface-to-air missiles (SAMs) are targeted early. This run, therefore, makes SAMs more attractive as a target to the aircraft. The results should indicate a preference for the SAMs, which in the baseline run is the lowest priority target. The focus of this research is to test the agent-based simulation's ability to execute a centralized grand strategy through a decentralized construct; this run explicitly tests this function.




Figure 4-4 Prefer SAMs – Targets Killed Over Time

TABLE 4-4 PREFER SAMs - WEAPONS EXPENDED

Prefer SAMs	664 total weapons				
	AGM-65	TMD-WC-SFW	JDAM	Mk-82	TOTAL
F15E	67	0	0	139	206
F16	46	0	0	104	150
F18EF	0	0	95	116	211
B-2	0	97	0	0	97
TOTAL	113	97	95	359	664

Note that the weapons expended are roughly the same for the “Prefer SAMs” scenario as compared to the baseline run. However, for the early part of the campaign, the Prefer SAMs scenario shows more targets killed per over time than in the baseline run. A review of the configuration file shows that more SAMs have been placed on the landscape than most other target types; 24% of targets are SAMs, while only 5% of the target base is comprised of C2 nodes. In the baseline run, the C2 nodes are higher priority, and therefore the agents prefer targeting these over SAMs. In this Prefer SAMs run, omit the agents kill more targets faster, which indicates that they may be targeting a more prolific target set than in the baseline run. For this comparison, an additional chart is needed in order to clearly see the desired change in preference.




Figure 4-5 Comparing SAMs targeted between Baseline and Prefer SAMs run

In Figure 4-5, the SAMs are targeted as a percentage of the overall target set. For the baseline run, shown in a dashed line, SAMs account for 54 of the total 250 targets attacked. For the Prefer SAMs run, shown in a solid line, the total is 69 SAMs. In comparison (not shown on chart), 22 C2 nodes are attacked in the baseline run, while only 14 are targeted in the Prefer SAMs run. In the baseline run, the C2 nodes are the highest

priority target. It is clear that the agents “prefer” the SAMs over C2 nodes when the sugar value is changed for the two target types. When the grand strategy is altered to indicate a preference for SAMs over C2 nodes, that strategy is reflected in the model when the agent rules are modified.

Over time in an air campaign, the Joint Force Air Component Commander will prefer certain targets types based on the phase of the campaign. In the early stages, air defense and leading edge armor may be preferred, while in the later stages of a campaign, infrastructure targets may be preferred. This grand strategy of phasing can be reflected in an agent-based simulation of an air campaign.

Further use of agent-based simulation in this field may explore the question of resource optimization. While not addressed explicitly in this work, questions remain: Is a self-organizing air campaign inherently more efficient in the use of scarce resources against a target set that must be attacked in phases? Can this modeling approach be used to test future hypotheses along these lines? If a system is made more adaptable and flexible when control is decentralized, and tactical (local) decisions follow a set of rules that reflect a grand strategy, perhaps this construct of a self-organizing air campaign will lead to a more efficient, locally-controlled but centrally commanded, use of airpower in major conflicts.

5. SUMMARY

This research is an existence proof to study the feasibility of a self-organizing air campaign for Joint air operations. Using agent-based modeling techniques, a simulation is constructed to test the concept of instituting a self-organizing system, wherein constructs of pilots, aircraft, weapons, escorts (all acting as agents) bid on fixed/known, as well as mobile/emergent, targets in a Joint air campaign. Through this piece, a body of work is begun that may lead to a more effective integration of weapons platforms (manned and unmanned aircraft) in Joint air operations. There are strong and numerous caveats to this conclusion, including the political hazards of allowing machines to acquire and destroy targets. Beyond this, the simple feasibility of a modeling technique is by no means sufficient standard for altering Air Force procurement policy or command and control strategy. In this research, the door has been opened to further research, proving only that the decentralized execution of a grand strategy is possible using agent-based simulation techniques, and that this approach may prove profitable for the prosecution of future air campaigns. As the U.S. moves towards a proliferation of target intelligence on the battlefield, and also towards the development of an Unmanned Combat Aerial Vehicle (UCAV) – this research indicates a possible use for the manned or unmanned platform in a future battlefield where target information is known quickly to all U.S. attack assets.

Deleted: n

The U.S. Air Force has been seeking to centralize the use of airpower in conflict, specifically seeking to provide a coherent and efficient use of airpower resources in support of the Commander's objectives in a joint conflict. This reflects an historic military dictum, which is to achieve unity of command over one's forces in order to make the most effective use of the forces. Winnefeld & Johnson raise the question regarding this striving towards unity of command, arguing that what is embedded in the term is unity of command and unity of control. While Air Force research and development efforts (DARPA) are aimed at increasing the analytic processing power at the level of the JFACC in order to better equip

the JFACC to manage a dynamic battle situation, it may be feasible to relieve the JFACC from the tactical battle management that is reflected in the ATO and MAAP – the assignment of specific aircraft to specific targets. Building on the observation of Winnefeld & Johnson, the study of a self-organizing air campaign is proposed. This reflects an acknowledgement for unity of command – here described as apportionment of effort over time against certain target classes, while arguing that unity of control can be accomplished procedurally, without burdensome target planning.

Because the argument can be made that the establishment of procedures for aircraft may be feasible, and provide for a reasonable air campaign simulation, the use of agent-based simulation methods is used to test this hypothesis:

It is possible to structure a feasible non-hierarchical approach to air tasking in the conduct of Joint air operations.

These methods enable the creation of individually-motivated autonomous agents on a landscape, each with embedded rules governing their behavior and objectives. The modeling environment from the Santa Fe Institute, Swarm, along with code built on top of the Swarm libraries, and Sugarscape is used in order to make use of leading implementations of agent-based simulations.

The Sugarscape rules are modified to allow for agent communication and coordination, and create a model, ABACUS, that features aircraft carrying weapons of various kill effectiveness against various targets classes. These aircraft coordinate while in flight, so that an aircraft with the most appropriate weapon against a particular target is “selected” from among the aircraft in proximity to the target. Without prior planning or tasking, the aircraft move about the battlefield, selecting targets and preferring those for whom the JFACC has expressed a high priority. When two aircraft converge on the same target, they communicate to determine the most appropriate course of action, and the “losing” aircraft breaks off its attack, seeking alternate targets.

It has been shown that a self-organizing air campaign is a feasible construct, at least in theory, and should inform design considerations for future JFACC research and development. This approach makes use of the distributed information available regarding

targets, and may possibly enable an increased use for unmanned combat air vehicles (UCAV). This research has not addressed the relative efficiency of a self-organizing approach. In order to do so, more complexity would have to be introduced into the model, in order to more closely abstract the combat environment. ABACUS could be extended to model topography, weather, target defenses and much more. The introduction of forces, natural and intentional, which work to limit the effectiveness of the aircraft would be a welcome next step to this research.

BIBLIOGRAPHY

BIBLIOGRAPHY

- Axtell, R. (2000). Why Agents? On the Varied Motivations for Agent Computing in the Social Sciences. Center on Social and Economic Dynamics Working Paper No. 17. Washington, DC, Brookings Institution.....
- Builder, C. H., S. C. Bankes, et al. (1999). Command Concepts: A Theory Derived from the Practice of Command and Control. Santa Monica, CA, RAND.
- Caddell, C. W. (2001). Balkan Traffic Cops. Airman Magazine.
- Carpenter, P. M. (1994). Joint Operations in the Gulf War: An Allison Analysis. School of Advanced Airpower Studies. Maxwell Air Force Base, Alabama, Air University.
- Coase, R. H. (1937). The Nature of the Firm (1937). The Nature of the Firm: Origins, Evolution and Development. O. E. Williamson and S. G. Winter. New York, Oxford University Press: 18-33.
- Cruetz, M. (1996). "Cellular Automata and Self Organized Criticality."
- Daniels, M. (1999). Integrating Simulation Technologies with Swarm. Agent Simulation: Applications, Models, Tools, University of Chicago.
- Davidow, W. H. and M. S. Malone (1992). The Virtual Corporation: Lessons from the World's Most Advanced Companies. New York, NY, HarperBusiness.
- Epstein, J. M. and R. L. Axtell (1996). Growing Artificial Societies from the Bottom Up (Complex Adaptive Systems). Washington DC, Brookings Institution Press.
- Friedman, G. and M. Friedman (1996). The Future of War: Power, Technology, and American World Dominance in the 21st Century. New York, NY, Crown Publishers.
- Frisbee, J. L. (1990). The Lessons of North Africa. Air Force Magazine. 73.
- Fukuyama, F. (1999). The Great Disruption: Human Nature and the Reconstitution of Social Order. New York, NY, The Free Press.
- Fukuyama, F. and S. Haeckel (1999). Leadership, Freedom, and Order in a Complex Adaptive System. Embracing Complexity 4: The Adaptive Enterprise, Cambridge, MA.
- Futrell, R. F. (1989). Ideas, Concepts, Doctrine: Basic Thinking in the United States Air Force 1907-1960. Maxwell AFB, AL, Air University Press.
- Futrell, R. F. (1988). "Tactical Employment of Strategic Air Power in Korea." Aerospace Power Journal Winter 1988.
- Futrell, R. F. (1983). The United States Air Force in Korea. Washington, D.C., Office of Air Force History.
- Gilbert, N. and K. G. Troitzsch (1999). Simulation for the Social Scientist. Philadelphia, PA, Open University Press.

- Heise, S. A. and H. S. Morse The DARPA JFACC Program: Modeling and Control of Military Operations. Arlington, VA,
- Holland, J. H. (1995). Hidden Order: How Adaptation Builds Complexity. Cambridge, MA, Perseus Books.
- Huber, A. F., P. S. Sauer, et al. (1996). The Virtual Air Combat Air Staff: The Promise of Information Technologies. Santa Monica, CA, RAND.
- Kropf, R. F. (1990). "The U.S. Air Force in Korea: Problems that Hindered the Effectiveness of Air Power." Airpower Journal **Spring 1990**.
- Loeb, V. (2001). Technology Changes Air War Tactics. Washington Post. Washington, DC: A16.
- Luttwak, E. N. (1992). Air Power in US Military Strategy. The Future of Air Power in the Aftermath of the Gulf War. R. H. Shultz, Jr. and J. Robert L. Pfaltzgraff. Maxwell Air Force Base, AL, Air University Press.
- Maes, P. (2000). Modeling Adaptive Autonomous Agents. Artificial Life: An Overview. C. Langton. Cambridge, MA, MIT Press.
- Martín, F. J., E. Plaza, et al. "Knowledge and Experience Reuse through Communication among Competent (Peer) Agents." International Journal of Software Engineering and Knowledge Engineering **9(3)**: 319-341.
- McEneaney, W. M., K. Ito, et al. (2001). NCSU Team: JFACC Final Report. Raleigh, NC, Department of Mathematics, North Carolina State University.
- McFadzean, D., D. Stewart, et al. (2000). A Computational Laboratory for Evolutionary Trade Networks, Iowa State University.....
- Millis, W., Ed. (1951). The Forrestal Diaries. New York, Viking.
- Minar, N., R. Burkhart, et al. (1996). "The Swarm Simulation System: A Toolkit for Building Multi-Agent Simulations." Available as: <http://www.santafe.edu/projects/swarm>
- Murray, W. (1992). Air Power since World War II: Consistent with Doctrine? The Future of Air Power in the Aftermath of the Gulf War. R. H. Shultz, Jr. and J. Robert L. Pfaltzgraff. Maxwell Air Force Base, AL, Air University Press.
- National Defense Panel. (1997). Transforming Defense: National Security in the 21st Century., National Defense Panel.
- Parunak, H. V. D. (1997). "Go to the Ant: Engineering Principles from Natural Multi-Agent Systems." Annals of Operations Research (Special Issue on Artificial Intelligence and Management Science) **75**(69-101).
- Penzias, A. (1995). Harmony: Business, Technology, & Life After Paperwork. New York, HarperBusiness.
- Rocha, L. M. (1999). From Artificial Life to Semiotic Agent Models: Review and Research Directions. Los Alamos, NM, Los Alamos National Laboratory.

- Schelling, T. C. (1978). Micromotives and Macrobehavior. New York, NY, W.W. Norton & Company.
- Sessions, S. D. and C. Jones (1993). Interoperability: A Desert Storm Case Study. Washington, DC, Institute for National Strategic Studies, National Defense University.
- Shen, W. and D. H. Norrie (1999). "Agent-Based Systems for Intelligent Manufacturing: A State-of-the-Art Survey." Knowledge and Information Systems, an International Journal 1(2): 129-156.
- Taylor, C. and D. Jefferson (1995). Artificial Life as a Tool for Biological Inquiry. Artificial Life: An Overview. C. G. Langton. Cambridge, MA, The MIT Press.
- Terna, P. (2001). "Creating Artificial Worlds: A Note on Sugarscape and Two Comments." Journal of Artificial Societies and Social Simulation 4(2).
- Tilford, E. H. (1991). Setup: What the Air Force Did in Vietnam and Why. Maxwell Air Force Base, AL, Air University Press.
- Waldrop, M. M. (1992). Complexity: The Emerging Science at the Edge of Order and Chaos. New York, Simon and Schuster.
- Warden, C. J. A. I., USAF (1989). The Air Campaign: Planning for Combat. Washington, DC, Pergamon-Brassey's International Defense Publishers.
- Winnefeld, J. A. and D. J. Johnson (1993). Joint Air Operations: Pursuit of Unity in Command and Control, 1942-1991. Annapolis, MD, Naval Institute Press.
- Womack, J. P., D. T. Jones, et al. (1991). The Machine that Changed the World: The Story of Lean Production, Harpercollins.
- Wooldridge, M. and N. R. Jennings (1995). "Intelligent Agents: Theory and Practice." Submitted to Knowledge Engineering Review.

APPENDIX A - ANNOTATED ABACUS CONFIGURATION FILE

Annotated ABACUS Configuration File

In this appendix, we will describe a configuration file for use within the ABACUS model. All values may be modified, and the configuration file must be saved to the appropriate directory as “abacus.cfg” in order to run the scenario featuring the modified values. For this section, the Courier New typeset is used to indicate text from the configuration file, while normal font indicates explanatory text.

In the initial section, the world is described in terms of total agents for the model run, total targets, how often a new aircraft is (generated or launched), how often a new target “pops up” or becomes visible to the aircraft, and the total simulation size for the theater (each square represents .25 mile square). In addition, certain display settings are found in this first section. Show agent choices refers to the appearance of a line on the display from the aircraft to its target of choice, while ShowTargetSelected indicates which targets have been selected for attack. While useful in viewing the simulation as aircraft are bumped from targets, the display may become cluttered with this setting turned on. 1 indicates on, 0 indicates off.

```
NumberOfAgents 100
NumberOfTargets 1000
NewAgentEveryXTurns 10
NewTargetEveryYTurns 30
HorizontalSimulationSize 600
VerticalSimulationSize 600
ShowAgentChoices 1
ShowTargetsSelected 1
```

This next section refers to values for targets. The five rows represent the five allowable target classes as follows: SAM (surface-to-air missile), APC (armored personnel carrier), SPH (self-propelled howitzer), Tank, and C2_Node (command and control node, deemed an important fixed target and representing the enemy’s

ability to coordinate his forces). Four numbers follow each target type, and represent, in order: the percentage of total targets over a simulation run that will be of this type, speed in terms of squares moved per ‘turn’, and sugarlow and sugarhigh, which indicate minimum and maximum values for sugar for each target within this type. The actual sugar value for each target will be a random value within this defined range. Recall that higher sugar values represent a higher priority for that target based on JFACC guidance.

```
Targets_percentages_speed_sugarlow_sugarhigh:
SAM 24 1 5 10
APC 48 0 7 30
SPH 11 1 15 20
Tank 12 1 22 26
C2_Node 5 0 35 50
```

The weapons are described in this next section. A descriptive name is followed by five rows. These five rows correspond to the target classes, and the probability of kill numbers for each target type are enumerated, along with an appropriateness indicator. In this way, we have constructed a killer-victim scoreboard, where each weapon can indicate its effectiveness and appropriateness to each target class. For example, the value 62 in the first row following the AGM-65 record indicates a 62% probability of kill for the AGM-65 against the first target class enumerated above, in this case, SAMs. Next to each weapon type is an indicator, “cheap” or “expensive.” While not currently used, this may be used later to indicate whether it is preferable to expend this munition on a target of low value, or return to base with the weapon undelivered.

```
Weapons_0=CheapOr1=Expensive_pk_Appropriateness_for_eac
h_target:
AGM-65 0
62 1
54 0
60 2
54 1
75 5
TMD-WC-SFW 1
50 0
```

```

90 0
95 2
99 5
10 1
JDAM 0
45 2
30 1
39 0
30 0
80 5
Mk-82 0
44 1
23 0
34 1
23 2
35 0

```

In the following section, we describe aircraft classes in terms of speed and time aloft, as well as target choices in order of priority, both while fully fueled, and while low on fuel. The thinking here is that the calculus changes as the aircraft becomes low on fuel. For example, it may be a better option to expend ordnance on a less-‘sweet’ target than to return to base without having expended ordnance. This may not be the case with expensive weapons, such as the Joint Direct Attack Munition or Tactical Munitions Dispenser.

The aircraft name is followed by the percentage of this aircraft type across the run of the simulation (this parallels the percentage of target types, noted in a previous section of the configuration file). Speed is denoted in terms of units or cells traversed during a ‘turn.’ Turns the aircraft remains on base while refueling (allows for variable sortie rates by aircraft class) is followed by the turns in air – the measure of time on station for an attacking aircraft, before it must abandon the attack and head towards base.

As discussed in the functional description of ABACUS, the next section details the priority for deciding targets. (The target priority is set by the assignment of sugar values according to target type, this value determines the relative value of sugar to the attacker’s decision to engage a specific target.) For the F-15E in this

example, the choice is the highest-priority target in the target list, followed by the closest among equally-high priority targets, followed by those for which it has the greatest chance of destruction among those which are equally-high priority, and which are equidistant from the aircraft's current location. Note the "1200" figure, which indicates the fuel level at which the aircraft calculus changes to reflect highest Pk, then closest, then sweetest from among candidate targets.

```
Agents_percentages_speed_TurnsToRefuel_TurnsInAir_Burnr
atePlaceholder_weapon_number_of_that_weapon:
F15E 30 4 3000 2500 1 AGM-65 2 Mk-82 4 .
Sweetest  Closest  HighestPK  1200  HighestPK  Closest
Sweetest
F16 30 3 4500 2500 1 Mk-82 6 AGM-65 2 .
Sweetest  Closest  HighestPK  1200  HighestPK  Closest
Sweetest
F18EF 30 3 4500 2500 1 JDAM 4 Mk-82 6 .
Sweetest  HighestPK  Closest  1200  HighestPK  Closest
Sweetest
B-2 10 2 13500 4883 1 TMD-WC-SFW 30 .
Sweetest  HighestPK  Closest  2440  Sweetest  Closest
HighestPK
```

The final section is used to arbitrate between competing agents.

```
AgentsRulesForCompetition:
Sugar > 5 HighestPK Closest BestFit
Sugar > 10 HighestPK Closest BestFit
Sugar > 15 Closest HighestPK BestFit

:
```


CURRICULUM VITAE

John Bordeaux was born on March 8, 1959 in Rockville Centre, New York, and is an American citizen. He graduated from Oceanside Senior High School, Oceanside, New York, in 1976. He received his Bachelor of Science from Christopher Newport University in 1989, and his Master of Science from George Mason University in 1998. After serving eight years in the United States Air Force, he was employed by The RAND Corporation from 1990 through 1998. From 1998 through 2000, he was head of software development, and then Knowledge Manager for the Modeling and Simulation Information Analysis Center, affiliated with the I.I.T. Research Institute. In 2001, he joined SRA International, Inc., in the position of Director, Knowledge Management. He received his Doctor of Philosophy in Public Policy with a concentration in Organizational Informatics from George Mason University in January, 2003.